



US009448803B2

(12) **United States Patent**  
**Lindholm et al.**

(10) **Patent No.:** **US 9,448,803 B2**  
(45) **Date of Patent:** **Sep. 20, 2016**

(54) **SYSTEM AND METHOD FOR HARDWARE SCHEDULING OF CONDITIONAL BARRIERS AND IMPATIENT BARRIERS**

(56) **References Cited**

U.S. PATENT DOCUMENTS

(71) Applicant: **NVIDIA Corporation**, Santa Clara, CA (US)

(72) Inventors: **John Erik Lindholm**, Saratoga, CA (US); **Tero Tapani Karras**, Helsinki (FI); **Timo Oskari Aila**, Tuusula (FI); **Samuli Matias Laine**, Vantaa (FI)

(73) Assignee: **NVIDIA Corporation**, Santa Clara, CA (US)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 646 days.

(21) Appl. No.: **13/794,578**

(22) Filed: **Mar. 11, 2013**

(65) **Prior Publication Data**

US 2014/0258693 A1 Sep. 11, 2014

(51) **Int. Cl.**

**G06F 9/30** (2006.01)  
**G06F 9/38** (2006.01)  
**G06F 9/48** (2006.01)  
**G06F 9/52** (2006.01)

(52) **U.S. Cl.**

CPC ..... **G06F 9/38** (2013.01); **G06F 9/30072** (2013.01); **G06F 9/30087** (2013.01); **G06F 9/3834** (2013.01); **G06F 9/3851** (2013.01); **G06F 9/4881** (2013.01); **G06F 9/522** (2013.01)

(58) **Field of Classification Search**

None

See application file for complete search history.

7,669,203 B2 2/2010 Samra et al.  
8,209,702 B1 6/2012 Roytman et al.  
2007/0113232 A1\* 5/2007 Collard ..... G06F 9/52 718/101  
2008/0256330 A1\* 10/2008 Wang ..... G06F 8/447 712/24  
2010/0153761 A1\* 6/2010 Nishioka ..... G06F 1/3203 713/323  
2011/0078417 A1\* 3/2011 Fahs ..... G06F 9/30145 712/216  
2012/0158886 A1\* 6/2012 O'Connell ..... H04L 67/22 709/217  
2014/0123150 A1\* 5/2014 Lindholm ..... G06F 9/4881 718/103  
2014/0282566 A1 9/2014 Lindholm et al.

FOREIGN PATENT DOCUMENTS

TW 201203111 A 1/2012

OTHER PUBLICATIONS

Office Action from Taiwanese Patent Application No. 102142872, dated Nov. 18, 2015.

(Continued)

*Primary Examiner* — William B Partridge

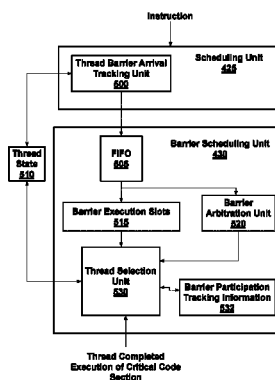
(74) *Attorney, Agent, or Firm* — Zilka-Kotab, PC

(57)

**ABSTRACT**

A method and a system are provided for hardware scheduling of barrier instructions. Execution of a plurality of threads to process instructions of a program that includes a barrier instruction is initiated, and when each thread reaches the barrier instruction during execution of program, it is determined whether the thread participates in the barrier instruction. The threads that participate in the barrier instruction are then serially executed to process one or more instructions of the program that follow the barrier instruction. A method and system are also provided for impatient scheduling of barrier instructions. When a portion of the threads that is greater than a minimum number of threads and less than all of the threads in the plurality of threads reaches the barrier instruction each of the threads in the portion is serially executed to process one or more instructions of the program that follow the barrier instruction.

**17 Claims, 11 Drawing Sheets**



(56)

**References Cited**

**OTHER PUBLICATIONS**

Non-Final Office Action from U.S. Appl. No. 13/844,541, dated Nov. 3, 2015.

Notice of Allowance from U.S. Appl. No. 13/844,541, dated May 4, 2016.

\* cited by examiner

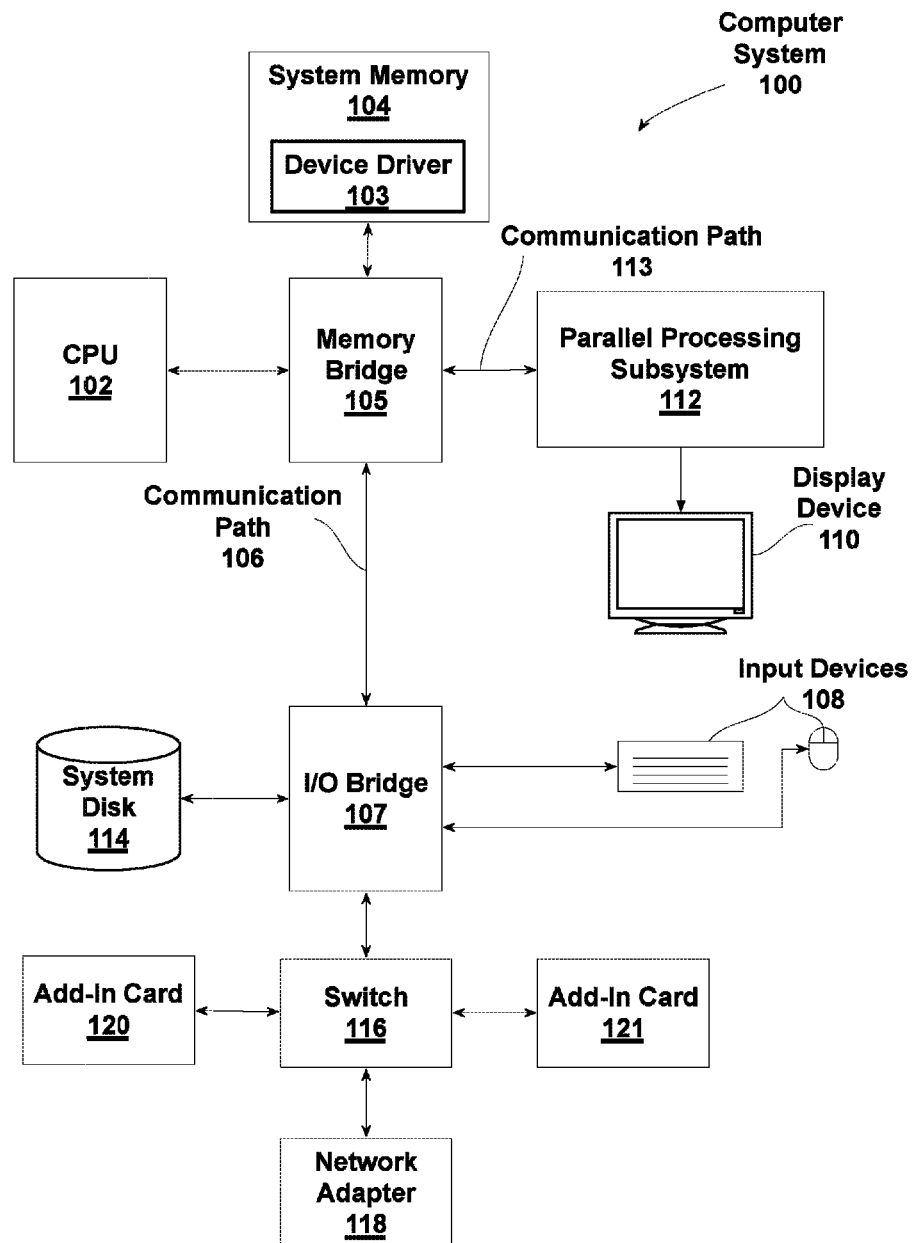


FIG. 1

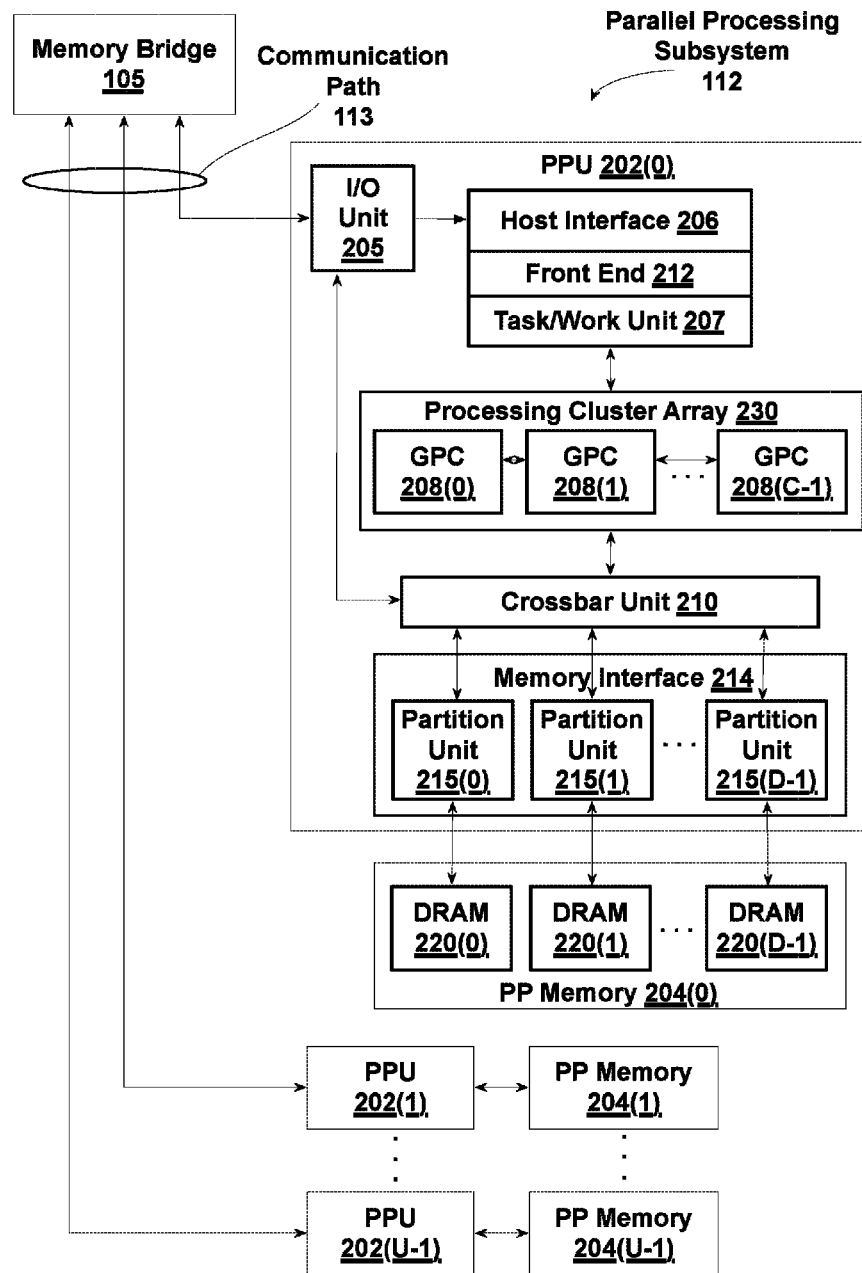


FIG. 2

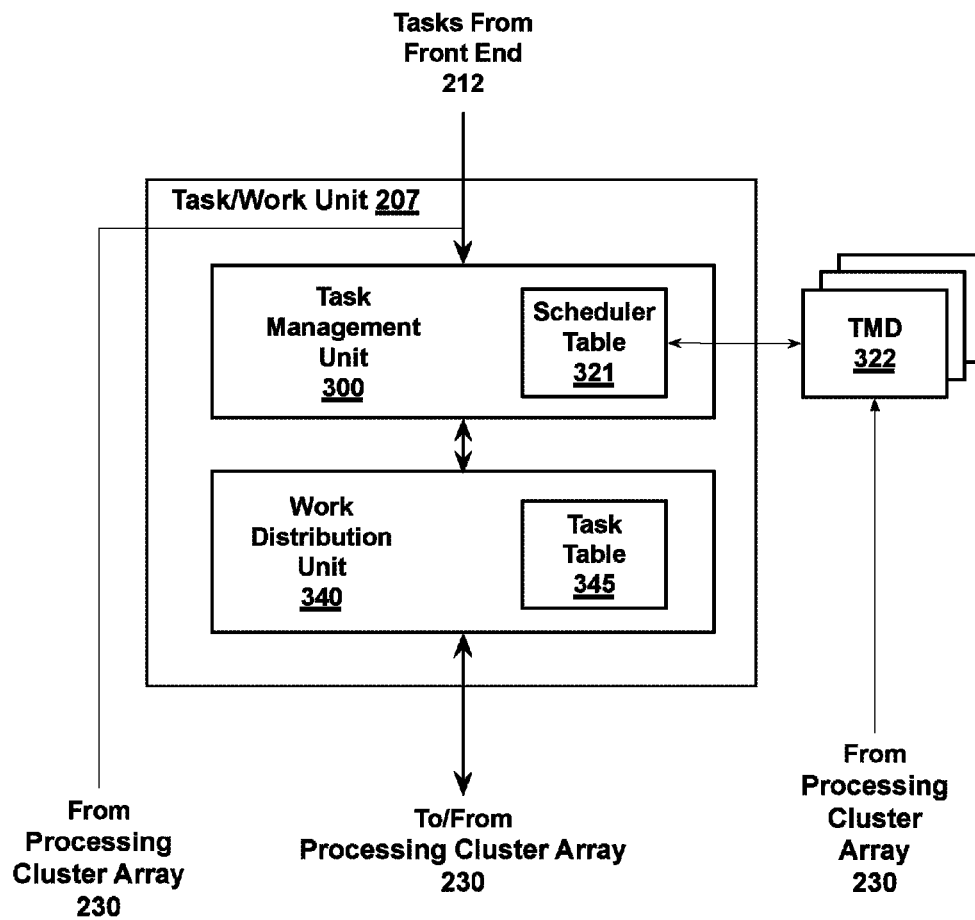


FIG. 3A

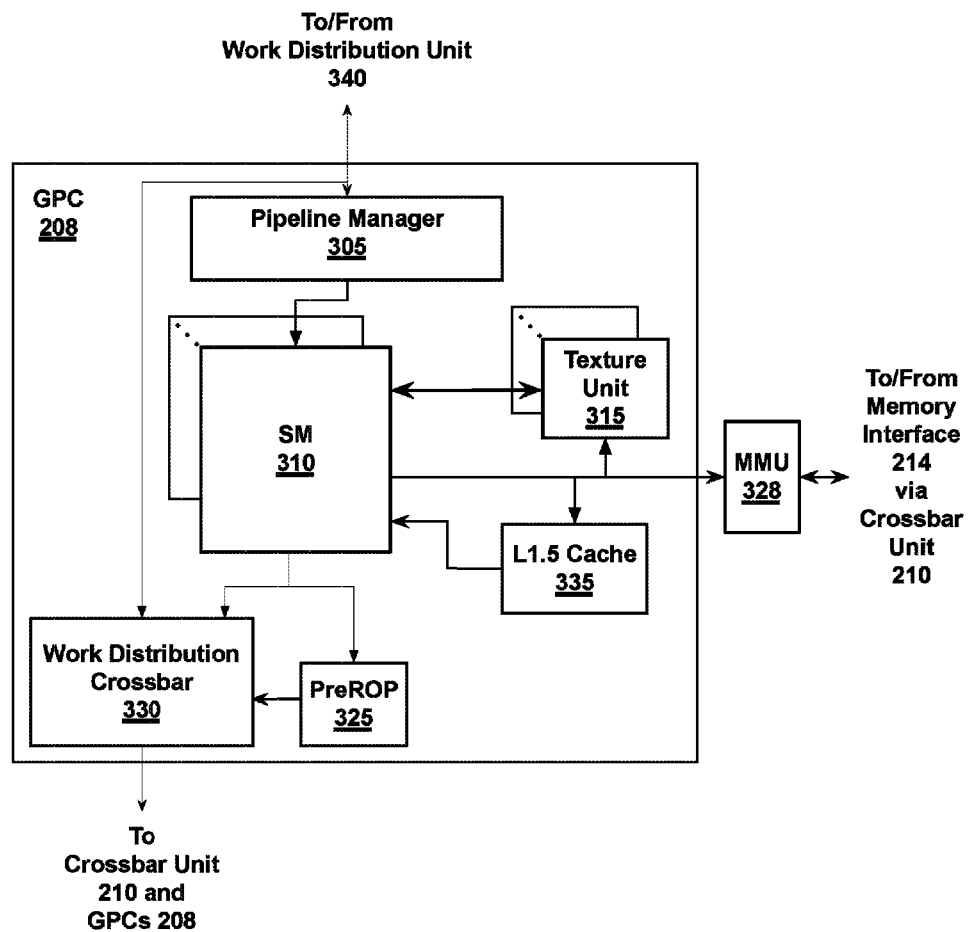


FIG. 3B

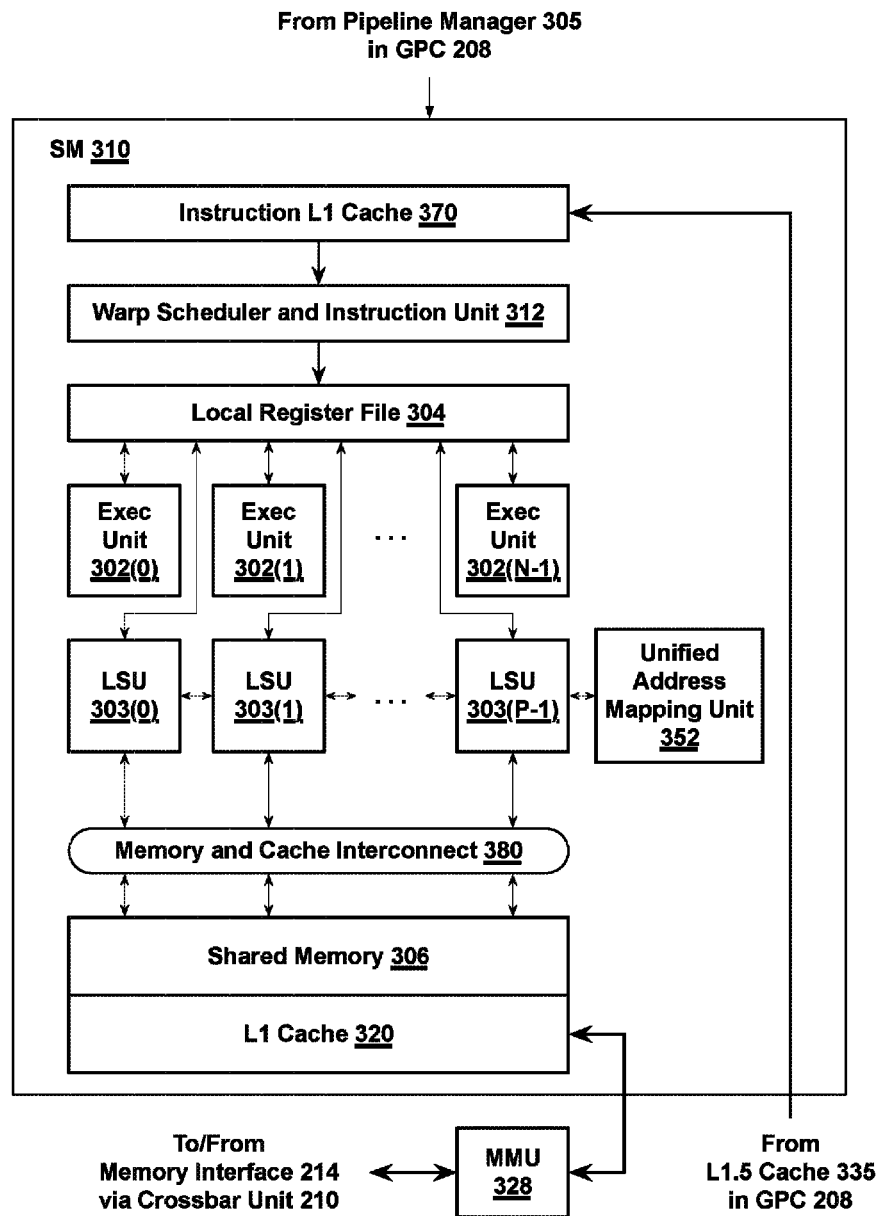


FIG. 3C

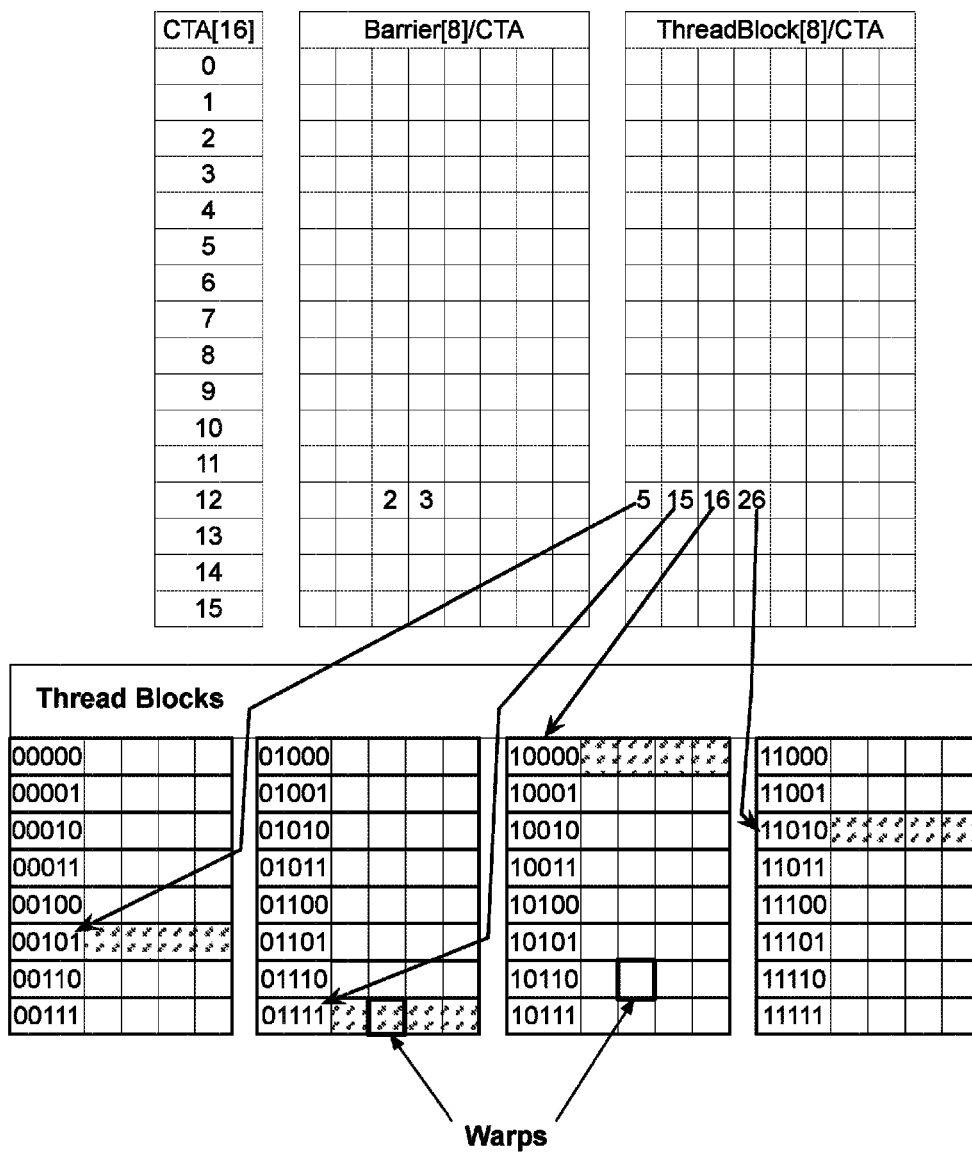


FIG. 4

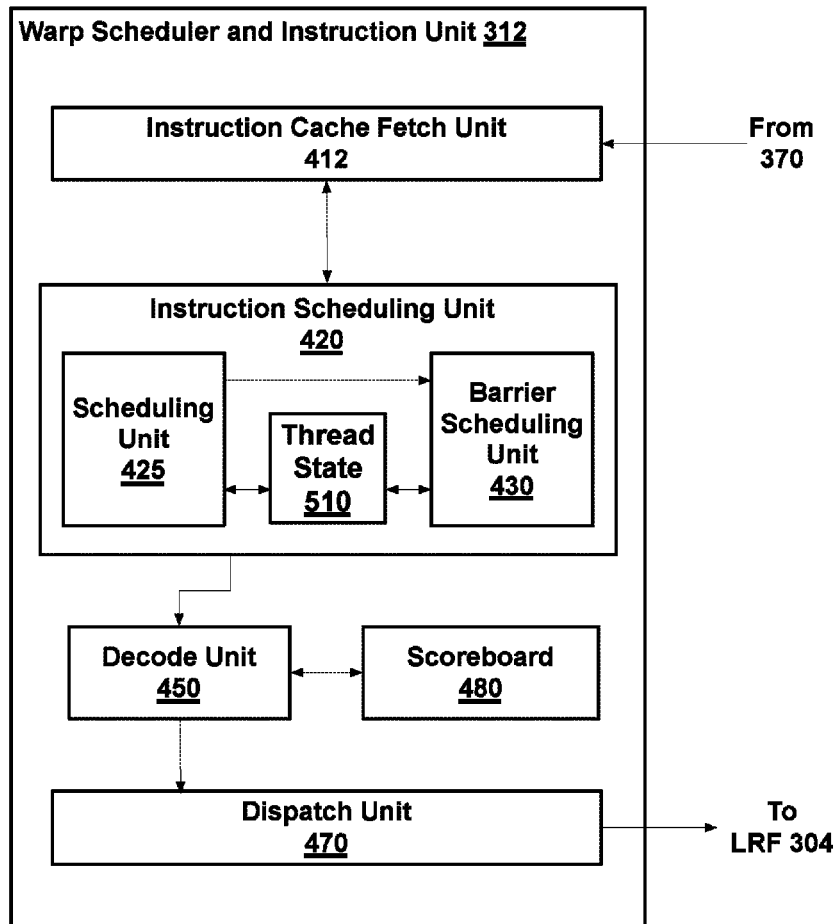


FIG. 5A

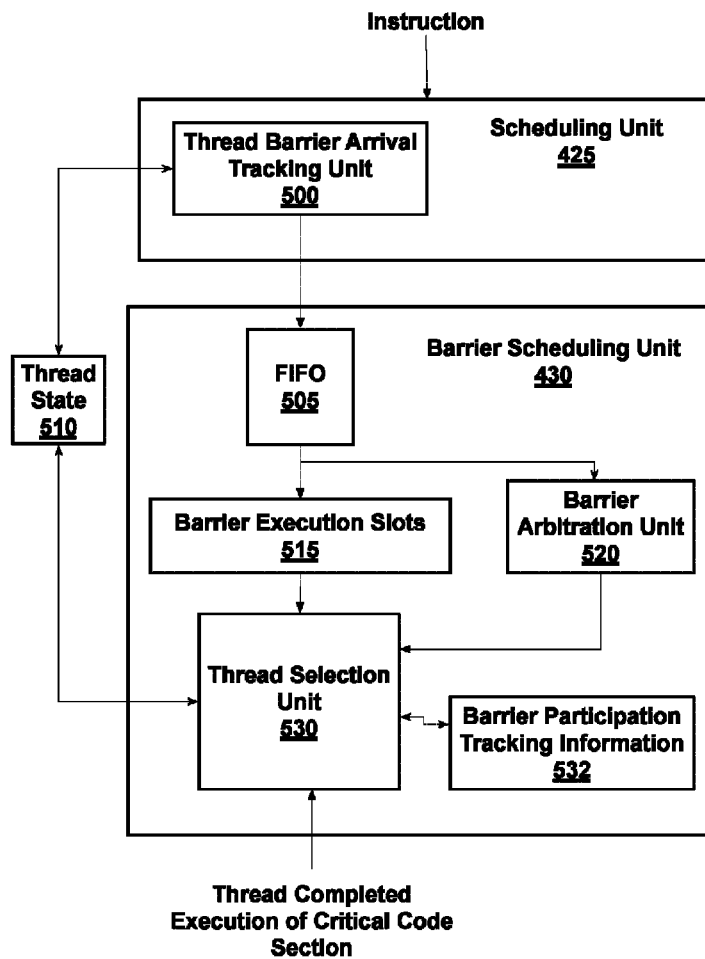


FIG. 5B

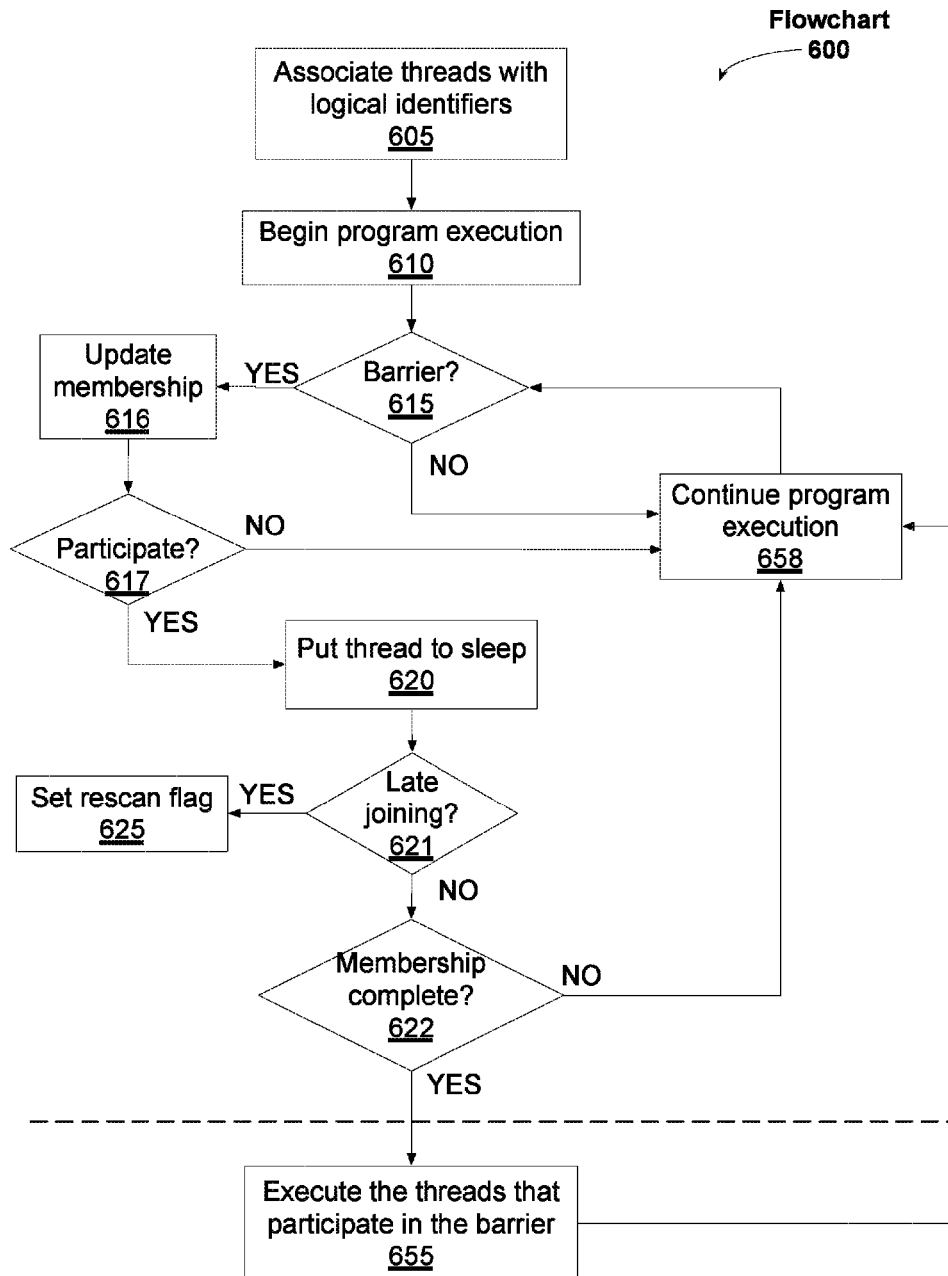


FIG. 6A

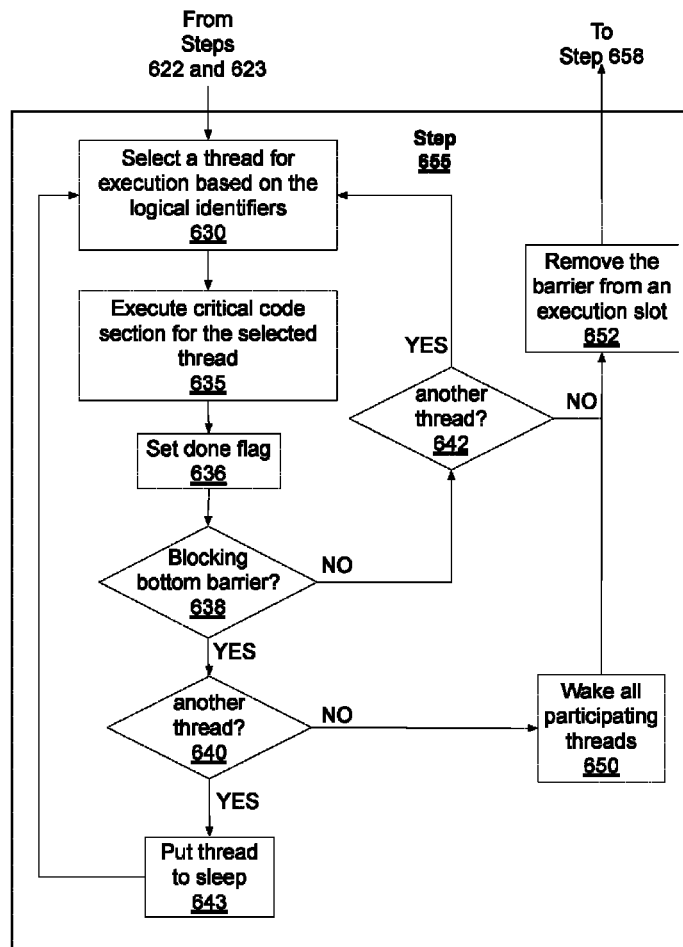


FIG. 6B

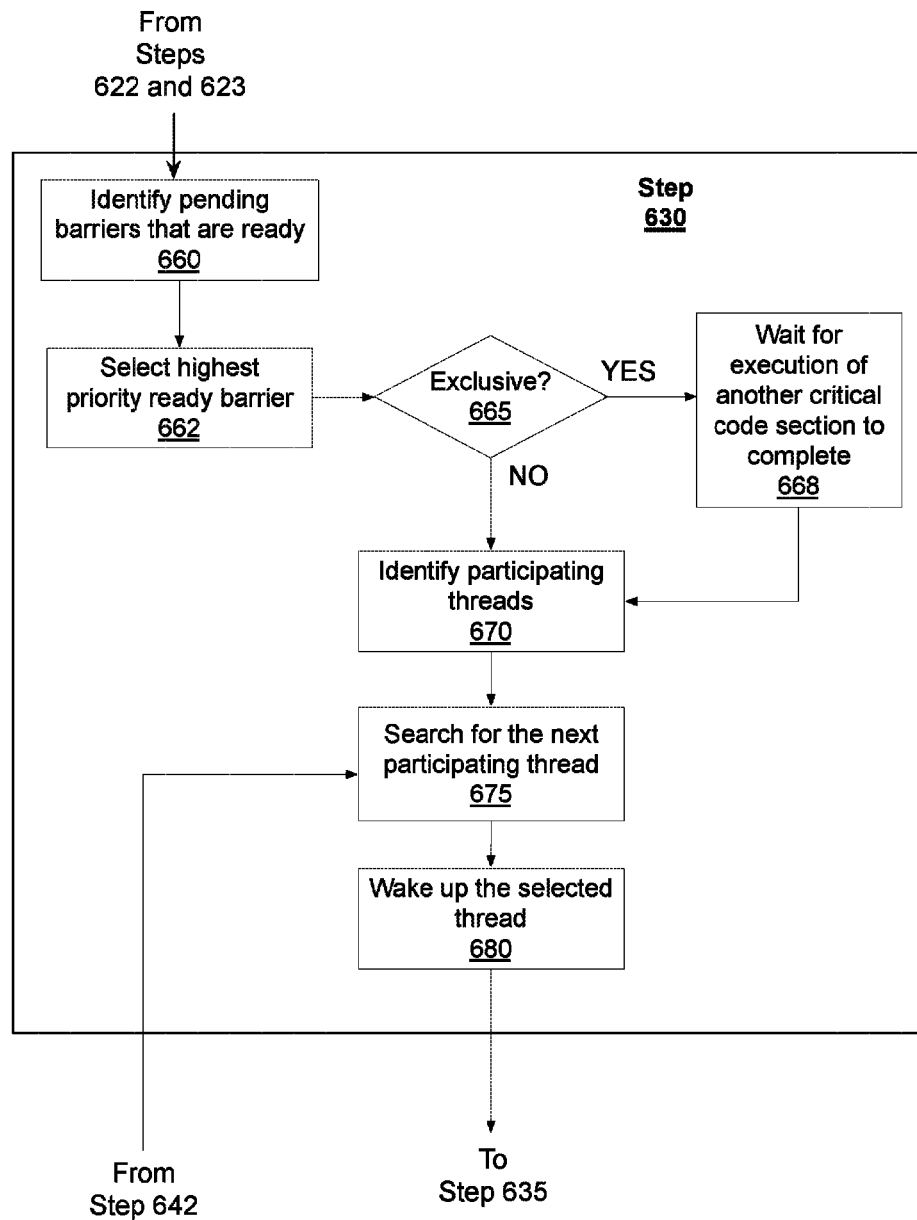


FIG. 6C

1

## SYSTEM AND METHOD FOR HARDWARE SCHEDULING OF CONDITIONAL BARRIERS AND IMPATIENT BARRIERS

This invention was made with Government support under LLNS subcontract B599861 awarded by DOE. The Government has certain rights in this invention.

### FIELD OF THE INVENTION

The present invention relates to program execution, and, more specifically to conditional or impatient barriers.

### BACKGROUND

Conventional parallel processing architectures support the execution of multiple threads. Particular operations that are performed during the execution of a program using a conventional parallel processing architecture may require synchronization of the multiple threads. Barrier instructions (or fence instructions) are used to synchronize the execution of multiple threads during execution of such a program. A scheduling unit within the parallel processing architecture recognizes the barrier instructions and ensures that all of the threads reach a particular barrier instruction before any of the threads executes an instruction subsequent to that particular barrier instruction. In some cases, it is not necessary to synchronize all of the threads at the barrier instruction and execution of the threads that do not require synchronization is unnecessarily delayed.

Thus, there is a need for addressing the issue of scheduling the execution threads to process barrier instructions and/or other issues associated with the prior art.

### SUMMARY

A system and method are provided for scheduling execution of conditional barriers. Execution of a plurality of threads to process instructions of a program that includes a barrier instruction is initiated, and when each thread reaches the barrier instruction during execution of program, it is determined whether the thread participates in the barrier instruction. The threads that participate in the barrier instruction are then serially executed to process one or more instructions of the program that follow the barrier instruction.

A method and system are also provided for impatient scheduling of barrier instructions. Execution of a plurality of threads to process instructions of a program that includes a barrier instruction is initiated, and when a portion of the threads that is greater than a minimum number of threads and less than all of the threads in the plurality of threads reaches the barrier instruction each of the threads in the portion is serially executed to process one or more instructions of the program that follow the barrier instruction.

### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram illustrating a computer system configured to implement one or more aspects of the present disclosure;

FIG. 2 is a block diagram of a parallel processing subsystem for the computer system of FIG. 1, according to one embodiment of the present disclosure;

FIG. 3A is a block diagram of the front end of FIG. 2, according to one embodiment of the present disclosure;

2

FIG. 3B is a block diagram of a general processing cluster within one of the parallel processing units of FIG. 2, according to one embodiment of the present disclosure;

FIG. 3C is a block diagram of a portion of the streaming multiprocessor of FIG. 3B, according to one embodiment of the present disclosure; and

FIG. 4 is a conceptual diagram illustrating thread blocks of a CTA, according to one example embodiment of the present disclosure;

FIG. 5A is a block diagram of the warp scheduler and instruction unit of FIG. 3C, according to one example embodiment of the present disclosure;

FIG. 5B is a block diagram of a portion of the scheduling unit and the barrier scheduling unit of FIG. 5A, according to one example embodiment of the present disclosure;

FIG. 6A illustrates a method for scheduling instructions for execution based on logical identifiers, according to one example embodiment of the present disclosure; and

FIG. 6B illustrates a method for performing a step shown in FIG. 6A, according to one example embodiment of the present disclosure.

FIG. 6C illustrates a method for performing a step shown in FIG. 6B, according to one example embodiment of the present disclosure.

### DETAILED DESCRIPTION

A system and method are provided for scheduling execution of conditional barriers that enable a subset of threads to participate in a barrier instruction. Execution of threads to process instructions of a program that includes a barrier instruction is initiated, and when each thread reaches the barrier instruction, it is determined whether the thread participates in the barrier instruction. The threads that participate in the barrier instruction are executed to process one or more instructions of the program that follow the barrier instruction. Threads that do not participate in the barrier instruction may continue execution without waiting for other threads to reach the barrier instruction.

More illustrative information will now be set forth regarding various optional architectures and features with which the foregoing technique may or may not be implemented, per the desires of the user. It should be strongly noted that the following information is set forth for illustrative purposes and should not be construed as limiting in any manner. Any of the following features may be optionally incorporated with or without the exclusion of other features described.

FIG. 1 is a block diagram illustrating a computer system 100 configured to implement one or more aspects of the present disclosure. Computer system 100 includes a central processing unit (CPU) 102 and a system memory 104 communicating via an interconnection path that may include a memory bridge 105. Memory bridge 105, which may be, e.g., a Northbridge chip, is connected via a bus or other communication path 106 (e.g., a HyperTransport link) to an I/O (input/output) bridge 107. I/O bridge 107, which may be, e.g., a Southbridge chip, receives user input from one or more user input devices 108 (e.g., keyboard, mouse) and forwards the input to CPU 102 via communication path 106 and memory bridge 105. A parallel processing subsystem 112 is coupled to memory bridge 105 via a bus or second communication path 113 (e.g., a Peripheral Component Interconnect (PCI) Express, Accelerated Graphics Port, or HyperTransport link); in one embodiment parallel processing subsystem 112 is a graphics subsystem that delivers pixels to a display device 110 (e.g., a conventional cathode

ray tube or liquid crystal display based monitor). A system disk **114** is also connected to I/O bridge **107**. A switch **116** provides connections between I/O bridge **107** and other components such as a network adapter **118** and various add-in cards **120** and **121**. Other components (not explicitly shown), including universal serial bus (USB) or other port connections, compact disc (CD) drives, digital video disc (DVD) drives, film recording devices, and the like, may also be connected to I/O bridge **107**. The various communication paths shown in FIG. 1, including the specifically named communications paths **106** and **113**, may be implemented using any suitable protocols, such as PCI Express, AGP (Accelerated Graphics Port), HyperTransport, or any other bus or point-to-point communication protocol(s), and connections between different devices may use different protocols as is known in the art.

In one embodiment, the parallel processing subsystem **112** incorporates circuitry optimized for graphics and video processing, including, for example, video output circuitry, and constitutes a graphics processing unit (GPU). In another embodiment, the parallel processing subsystem **112** incorporates circuitry optimized for general purpose processing, while preserving the underlying computational architecture, described in greater detail herein. In yet another embodiment, the parallel processing subsystem **112** may be integrated with one or more other system elements in a single subsystem, such as joining the memory bridge **105**, CPU **102**, and I/O bridge **107** to form a system on chip (SoC).

It will be appreciated that the system shown herein is illustrative and that variations and modifications are possible. The connection topology, including the number and arrangement of bridges, the number of CPUs **102**, and the number of parallel processing subsystems **112**, may be modified as desired. For instance, in some embodiments, system memory **104** is connected to CPU **102** directly rather than through a bridge, and other devices communicate with system memory **104** via memory bridge **105** and CPU **102**. In other alternative topologies, parallel processing subsystem **112** is connected to I/O bridge **107** or directly to CPU **102**, rather than to memory bridge **105**. In still other embodiments, I/O bridge **107** and memory bridge **105** might be integrated into a single chip instead of existing as one or more discrete devices. Large embodiments may include two or more CPUs **102** and two or more parallel processing systems **112**. The particular components shown herein are optional; for instance, any number of add-in cards or peripheral devices might be supported. In some embodiments, switch **116** is eliminated, and network adapter **118** and add-in cards **120**, **121** connect directly to I/O bridge **107**.

FIG. 2 illustrates a parallel processing subsystem **112**, according to one embodiment of the present disclosure. As shown, parallel processing subsystem **112** includes one or more parallel processing units (PPUs) **202**, each of which is coupled to a local parallel processing (PP) memory **204**. In general, a parallel processing subsystem includes a number  $U$  of PPUs, where  $U \geq 1$ . (Herein, multiple instances of like objects are denoted with reference numbers identifying the object and parenthetical numbers identifying the instance where needed.) PPUs **202** and parallel processing memories **204** may be implemented using one or more integrated circuit devices, such as programmable processors, application specific integrated circuits (ASICs), or memory devices, or in any other technically feasible fashion.

Referring again to FIG. 1 as well as FIG. 2, in some embodiments, some or all of PPUs **202** in parallel processing subsystem **112** are graphics processors with rendering pipelines that can be configured to perform various operations

related to generating pixel data from graphics data supplied by CPU **102** and/or system memory **104** via memory bridge **105** and the second communication path **113**, interacting with local parallel processing memory **204** (which can be used as graphics memory including, e.g., a conventional frame buffer) to store and update pixel data, delivering pixel data to display device **110**, and the like. In some embodiments, parallel processing subsystem **112** may include one or more PPUs **202** that operate as graphics processors and one or more other PPUs **202** that are used for general-purpose computations. The PPUs may be identical or different, and each PPU may have a dedicated parallel processing memory device(s) or no dedicated parallel processing memory device(s). One or more PPUs **202** in parallel processing subsystem **112** may output data to display device **110** or each PPU **202** in parallel processing subsystem **112** may output data to one or more display devices **110**.

In operation, CPU **102** is the master processor of computer system **100**, controlling and coordinating operations of other system components. In particular, CPU **102** issues commands that control the operation of PPUs **202**. In some embodiments, CPU **102** writes a stream of commands for each PPU **202** to a data structure (not explicitly shown in either FIG. 1 or FIG. 2) that may be located in system memory **104**, parallel processing memory **204**, or another storage location accessible to both CPU **102** and PPU **202**. A pointer to each data structure is written to a pushbuffer to initiate processing of the stream of commands in the data structure. The PPU **202** reads command streams from one or more pushbuffers and then executes commands asynchronously relative to the operation of CPU **102**. Execution priorities may be specified for each pushbuffer by an application program via the device driver **103** to control scheduling of the different pushbuffers.

Referring back now to FIG. 2 as well as FIG. 1, each PPU **202** includes an I/O (input/output) unit **205** that communicates with the rest of computer system **100** via communication path **113**, which connects to memory bridge **105** (or, in one alternative embodiment, directly to CPU **102**). The connection of PPU **202** to the rest of computer system **100** may also be varied. In some embodiments, parallel processing subsystem **112** is implemented as an add-in card that can be inserted into an expansion slot of computer system **100**. In other embodiments, a PPU **202** can be integrated on a single chip with a bus bridge, such as memory bridge **105** or I/O bridge **107**. In still other embodiments, some or all elements of PPU **202** may be integrated on a single chip with CPU **102**.

In one embodiment, communication path **113** is a PCI Express link, in which dedicated lanes are allocated to each PPU **202**, as is known in the art. Other communication paths may also be used. An I/O unit **205** generates packets (or other signals) for transmission on communication path **113** and also receives all incoming packets (or other signals) from communication path **113**, directing the incoming packets to appropriate components of PPU **202**. For example, commands related to processing tasks may be directed to a host interface **206**, while commands related to memory operations (e.g., reading from or writing to parallel processing memory **204**) may be directed to a memory crossbar unit **210**. Host interface **206** reads each pushbuffer and outputs the command stream stored in the pushbuffer to a front end **212**.

Each PPU **202** advantageously implements a highly parallel processing architecture. As shown in detail, PPU **202(0)** includes a processing cluster array **230** that includes a

number C of general processing clusters (GPCs) 208, where  $C \geq 1$ . Each GPC 208 is capable of executing a large number (e.g., hundreds or thousands) of threads concurrently, where each thread is an instance of a program. In various applications, different GPCs 208 may be allocated for processing different types of programs or for performing different types of computations. The allocation of GPCs 208 may vary dependent on the workload arising for each type of program or computation.

GPCs 208 receive processing tasks to be executed from a work distribution unit within a task/work unit 207. The work distribution unit receives pointers to processing tasks that are encoded as task metadata (TMD) and stored in memory. The pointers to TMDs are included in the command stream that is stored as a pushbuffer and received by the front end unit 212 from the host interface 206. Processing tasks that may be encoded as TMDs include indices of data to be processed, as well as state parameters and commands defining how the data is to be processed (e.g., what program is to be executed). The task/work unit 207 receives tasks from the front end 212 and ensures that GPCs 208 are configured to a valid state before the processing specified by each one of the TMDs is initiated. A priority may be specified for each TMD that is used to schedule execution of the processing task. Processing tasks can also be received from the processing cluster array 230. Optionally, the TMD can include a parameter that controls whether the TMD is added to the head or the tail for a list of processing tasks (or list of pointers to the processing tasks), thereby providing another level of control over priority.

Memory interface 214 includes a number D of partition units 215 that are each directly coupled to a portion of parallel processing memory 204, where  $D \geq 1$ . As shown, the number of partition units 215 generally equals the number of dynamic random access memory (DRAM) 220. In other embodiments, the number of partition units 215 may not equal the number of memory devices. Persons of ordinary skill in the art will appreciate that DRAM 220 may be replaced with other suitable storage devices and can be of generally conventional design. A detailed description is therefore omitted. Render targets, such as frame buffers or texture maps may be stored across DRAMs 220, allowing partition units 215 to write portions of each render target in parallel to efficiently use the available bandwidth of parallel processing memory 204.

Any one of GPCs 208 may process data to be written to any of the DRAMs 220 within parallel processing memory 204. Crossbar unit 210 is configured to route the output of each GPC 208 to the input of any partition unit 215 or to another GPC 208 for further processing. GPCs 208 communicate with memory interface 214 through crossbar unit 210 to read from or write to various external memory devices. In one embodiment, crossbar unit 210 has a connection to memory interface 214 to communicate with I/O unit 205, as well as a connection to local parallel processing memory 204, thereby enabling the processing cores within the different GPCs 208 to communicate with system memory 104 or other memory that is not local to PPU 202. In the embodiment shown in FIG. 2, crossbar unit 210 is directly connected with I/O unit 205. Crossbar unit 210 may use virtual channels to separate traffic streams between the GPCs 208 and partition units 215.

Again, GPCs 208 can be programmed to execute processing tasks relating to a wide variety of applications, including but not limited to, linear and nonlinear data transforms, filtering of video and/or audio data, modeling operations (e.g., applying laws of physics to determine position, veloc-

ity and other attributes of objects), image rendering operations (e.g., tessellation shader, vertex shader, geometry shader, and/or pixel shader programs), and so on. PPUs 202 may transfer data from system memory 104 and/or local parallel processing memories 204 into internal (on-chip) memory, process the data, and write result data back to system memory 104 and/or local parallel processing memories 204, where such data can be accessed by other system components, including CPU 102 or another parallel processing subsystem 112.

A PPU 202 may be provided with any amount of local parallel processing memory 204, including no local memory, and may use local memory and system memory in any combination. For instance, a PPU 202 can be a graphics processor in a unified memory architecture (UMA) embodiment. In such embodiments, little or no dedicated graphics (parallel processing) memory would be provided, and PPU 202 would use system memory exclusively or almost exclusively. In UMA embodiments, a PPU 202 may be integrated into a bridge chip or processor chip or provided as a discrete chip with a high-speed link (e.g., PCI Express) connecting the PPU 202 to system memory via a bridge chip or other communication means.

As noted above, any number of PPUs 202 can be included in a parallel processing subsystem 112. For instance, multiple PPUs 202 can be provided on a single add-in card, or multiple add-in cards can be connected to communication path 113, or one or more of PPUs 202 can be integrated into a bridge chip. PPUs 202 in a multi-PPU system may be identical to or different from one another. For instance, different PPUs 202 might have different numbers of processing cores, different amounts of local parallel processing memory, and so on. Where multiple PPUs 202 are present, those PPUs may be operated in parallel to process data at a higher throughput than is possible with a single PPU 202. Systems incorporating one or more PPUs 202 may be implemented in a variety of configurations and form factors, including desktop, laptop, or handheld personal computers, servers, workstations, game consoles, embedded systems, and the like.

#### Multiple Concurrent Task Scheduling

Multiple processing tasks may be executed concurrently on the GPCs 208 and a processing task may generate one or more "child" processing tasks during execution. The task/work unit 207 receives the tasks and dynamically schedules the processing tasks and child processing tasks for execution by the GPCs 208.

FIG. 3A is a block diagram of the task/work unit 207 of FIG. 2, according to one embodiment of the present disclosure. The task/work unit 207 includes a task management unit 300 and the work distribution unit 340. The task management unit 300 organizes tasks to be scheduled based on execution priority levels. For each priority level, the task management unit 300 stores a list of pointers to the TMDs 322 corresponding to the tasks in the scheduler table 321, where the list may be implemented as a linked list. The TMDs 322 may be stored in the PP memory 204 or system memory 104. The rate at which the task management unit 300 accepts tasks and stores the tasks in the scheduler table 321 is decoupled from the rate at which the task management unit 300 schedules tasks for execution. Therefore, the task management unit 300 may collect several tasks before scheduling the tasks. The collected tasks may then be scheduled based on priority information or using other techniques, such as round-robin scheduling.

The work distribution unit **340** includes a task table **345** with slots that may each be occupied by the TMD **322** for a task that is being executed. The task management unit **300** may schedule tasks for execution when there is a free slot in the task table **345**. When there is not a free slot, a higher priority task that does not occupy a slot may evict a lower priority task that does occupy a slot. When a task is evicted, the task is stopped, and if execution of the task is not complete, then a pointer to the task is added to a list of task pointers to be scheduled so that execution of the task will resume at a later time. When a child processing task is generated, during execution of a task, a pointer to the child task is added to the list of task pointers to be scheduled. A child task may be generated by a TMD **322** executing in the processing cluster array **230**.

Unlike a task that is received by the task/work unit **207** from the front end **212**, child tasks are received from the processing cluster array **230**. Child tasks are not inserted into pushbuffers or transmitted to the front end. The CPU **102** is not notified when a child task is generated or data for the child task is stored in memory. Another difference between the tasks that are provided through pushbuffers and child tasks is that the tasks provided through the pushbuffers are defined by the application program whereas the child tasks are dynamically generated during execution of the tasks.

#### Task Processing Overview

FIG. 3B is a block diagram of a GPC **208** within one of the PPUs **202** of FIG. 2, according to one embodiment of the present disclosure. Each GPC **208** may be configured to execute a large number of threads in parallel, where the term “thread” refers to an instance of a particular program executing on a particular set of input data. In some embodiments, single-instruction, multiple-data (SIMD) instruction issue techniques are used to support parallel execution of a large number of threads without providing multiple independent instruction units. In other embodiments, single-instruction, multiple-thread (SMT) techniques are used to support parallel execution of a large number of generally synchronized threads, using a common instruction unit configured to issue instructions to a set of processing engines within each one of the GPCs **208**. Unlike a SIMD execution regime, where all processing engines typically execute identical instructions, SMT execution allows different threads to more readily follow divergent execution paths through a given thread program. Persons of ordinary skill in the art will understand that a SIMD processing regime represents a functional subset of a SMT processing regime. Similarly, a SISD (single-instruction single-data) or a MIMD (multiple-instruction multiple-data) regime also represents a functional subset of a SMT processing regime.

Operation of GPC **208** is advantageously controlled via a pipeline manager **305** that distributes processing tasks to streaming multiprocessors (SMs) **310**. Pipeline manager **305** may also be configured to control a work distribution crossbar **330** by specifying destinations for processed data output by SMs **310**.

In one embodiment, each GPC **208** includes a number  $M$  of SMs **310**, where  $M \geq 1$ , each SM **310** configured to process one or more thread groups. Also, each SM **310** advantageously includes an identical set of functional execution units (e.g., execution units and load-store units-shown as Exec units **302** and LSUs **303** in FIG. 3C) that may be pipelined, allowing a new instruction to be issued before a previous instruction has finished, as is known in the art. Any combination of functional execution units may be provided.

In one embodiment, the functional units support a variety of operations including integer and floating point arithmetic (e.g., addition and multiplication), comparison operations, Boolean operations (AND, OR, XOR), bit-shifting, and computation of various algebraic functions (e.g., planar interpolation, trigonometric, exponential, and logarithmic functions, etc.); and the same functional unit hardware can be leveraged to perform different operations.

The series of instructions transmitted to a particular GPC **208** are executed by one or more threads, as previously defined herein, and the collection of a certain number of concurrently executing threads across the parallel processing engines (not shown) within an SM **310** is referred to herein as a “warp” or “thread group.” As used herein, a “thread group” refers to a group of one or more threads concurrently executing the same program on different input data, with one thread of the group being assigned to a different processing engine within an SM **310**. A thread group may include fewer threads than the number of processing engines within the SM **310**, in which case some processing engines will be idle during cycles when that thread group is being processed. A thread group may also include more threads than the number of processing engines within the SM **310**, in which case processing will take place over consecutive clock cycles. Since each SM **310** can support up to  $G$  thread groups concurrently, it follows that up to  $G \cdot M$  thread groups can be executing in GPC **208** at any given time.

Additionally, a plurality of related warps may be active (in different phases of execution) at the same time within an SM **310**. This collection of thread groups is referred to herein as a “cooperative thread array” (“CTA”) or “thread array.” Each CTA comprises a programmer-specified number of warps executing in the same SM **310**. One or more CTAs can potentially execute concurrently in the same SM **310**. The size of a CTA is generally determined by the programmer and the amount of hardware resources, such as memory or registers, available to the CTA.

Each SM **310** contains a level one (L1) cache (shown in FIG. 3C) or uses space in a corresponding L1 cache outside of the SM **310** that is used to perform load and store operations. Each SM **310** also has access to level two (L2) caches that are shared among all GPCs **208** and may be used to transfer data between threads. Finally, SMs **310** also have access to off-chip “global” memory, which can include, e.g., parallel processing memory **204** and/or system memory **104**. It is to be understood that any memory external to PPU **202** may be used as global memory. Additionally, a level one-point-five (L1.5) cache **335** may be included within the GPC **208**, configured to receive and hold data fetched from memory via memory interface **214** requested by SM **310**, including instructions, uniform data, and constant data, and provide the requested data to SM **310**. Embodiments having multiple SMs **310** in GPC **208** beneficially share common instructions and data cached in L1.5 cache **335**.

Each GPC **208** may include a memory management unit (MMU) **328** that is configured to map virtual addresses into physical addresses. In other embodiments, MMU(s) **328** may reside within the memory interface **214**. The MMU **328** includes a set of page table entries (PTEs) used to map a virtual address to a physical address. The MMU **328** may include address translation lookaside buffers (TLB) or caches which may reside within multiprocessor SM **310** or the L1 cache or GPC **208**. The physical address is processed to distribute data access locality to allow efficient request

interleaving among partition units **215**. The cache line index may be used to determine whether or not a request for a cache line is a hit or miss.

In graphics and computing applications, a GPC **208** may be configured such that each SM **310** is coupled to a texture unit **315** for performing texture mapping operations, e.g., determining texture sample positions, reading texture data, and filtering the texture data. Texture data is read from an internal texture L1 cache (not shown) or in some embodiments from the L1 cache within SM **310** and is fetched from an L2 cache that is shared between all GPCs **208**, parallel processing memory **204**, or system memory **104**, as needed. Each SM **310** outputs processed tasks to work distribution crossbar **330** in order to provide the processed task to another GPC **208** for further processing or to store the processed task in an L2 cache, parallel processing memory **204**, or system memory **104** via crossbar unit **210**. A preROP (pre-raster operations) **325** is configured to receive data from SM **310**, direct data to ROP units within partition units **215**, and perform optimizations for color blending, organize pixel color data, and perform address translations.

It will be appreciated that the core architecture described herein is illustrative and that variations and modifications are possible. Any number of processing units, e.g., SMs **310** or texture units **315**, preROPs **325** may be included within a GPC **208**. Further, as shown in FIG. **2**, a PPU **202** may include any number of GPCs **208** that are advantageously functionally similar to one another so that execution behavior does not depend on which GPC **208** receives a particular processing task. Further, each GPC **208** advantageously operates independently of other GPCs **208** using separate and distinct processing units, L1 caches to execute tasks for one or more application programs.

Persons of ordinary skill in the art will understand that the architecture described in FIGS. **1**, **2**, **3A**, and **3B** in no way limits the scope of the present invention and that the techniques taught herein may be implemented on any properly configured processing unit, including, without limitation, one or more CPUs, one or more multi-core CPUs, one or more PPUs **202**, one or more GPCs **208**, one or more graphics or special purpose processing units, or the like, without departing the scope of the present invention.

In embodiments of the present invention, it is desirable to use PPU **202** or other processor(s) of a computing system to execute general-purpose computations using thread arrays. Each thread in the thread array is assigned a unique thread identifier (“thread ID”) that is accessible to the thread during the thread’s execution. The thread ID, which can be defined as a one-dimensional or multi-dimensional numerical value controls various aspects of the thread’s processing behavior. For instance, a thread ID may be used to determine which portion of the input data set a thread is to process and/or to determine which portion of an output data set a thread is to produce or write.

A sequence of per-thread instructions may include at least one instruction that defines a cooperative behavior between the representative thread and one or more other threads of the thread array. For example, the sequence of per-thread instructions might include an instruction to suspend execution of operations for the representative thread at a particular point in the sequence until such time as one or more of the other threads reach that particular point, an instruction for the representative thread to store data in a shared memory to which one or more of the other threads have access, an instruction for the representative thread to atomically read and update data stored in a shared memory to which one or more of the other threads have access based on their thread

IDs, or the like. The program can also include an instruction to compute an address in the shared memory from which data is to be read, with the address being a function of thread ID. By defining suitable functions and providing synchronization techniques, data can be written to a given location in shared memory by one thread of a CTA and read from that location by a different thread of the same CTA in a predictable manner. Consequently, any desired pattern of data sharing among threads can be supported, and any thread in a CTA can share data with any other thread in the same CTA. The extent, if any, of data sharing among threads of a CTA is determined by the program; thus, it is to be understood that in a particular application that uses CTAs, the threads of a CTA might or might not actually share data with each other, depending on the program, and the terms “CTA” and “thread array” are used synonymously herein.

FIG. **3C** is a block diagram of the SM **310** of FIG. **3B**, according to one embodiment of the present disclosure. The SM **310** includes an instruction L1 cache **370** that is configured to receive instructions and constants from memory via L1.5 cache **335**. A warp scheduler and instruction unit **312** receives instructions and constants from the instruction L1 cache **370** and controls local register file **304** and SM **310** functional units according to the instructions and constants. The SM **310** functional units include N exec (execution or processing) units **302** and P load-store units (LSU) **303**.

SM **310** provides on-chip (internal) data storage with different levels of accessibility. Special registers (not shown) are readable but not writeable by LSU **303** and are used to store parameters defining each thread’s “position.” In one embodiment, special registers include one register per thread (or per exec unit **302** within SM **310**) that stores a thread ID; each thread ID register is accessible only by a respective one of the exec unit **302**. Special registers may also include additional registers, readable by all threads that execute the same processing task represented by a TMD **322** (or by all LSUs **303**) that store a CTA identifier, the CTA dimensions, the dimensions of a grid to which the CTA belongs (or queue position if the TMD **322** encodes a queue task instead of a grid task), and an identifier of the TMD **322** to which the CTA is assigned.

If the TMD **322** is a grid TMD, execution of the TMD **322** causes a fixed number of CTAs to be launched and executed to process the fixed amount of data stored in the queue. The number of CTAs is specified as the product of the grid width, height, and depth. The fixed amount of data may be stored in the TMD **322** or the TMD **322** may store a pointer to the data that will be processed by the CTAs. The TMD **322** also stores a starting address of the program that is executed by the CTAs.

If the TMD **322** is a queue TMD, then a queue feature of the TMD **322** is used, meaning that the amount of data to be processed is not necessarily fixed. Queue entries store data for processing by the CTAs assigned to the TMD **322**. The queue entries may also represent a child task that is generated by another TMD **322** during execution of a thread, thereby providing nested parallelism. Typically, execution of the thread, or CTA that includes the thread, is suspended until execution of the child task completes. The queue may be stored in the TMD **322** or separately from the TMD **322**, in which case the TMD **322** stores a queue pointer to the queue. Advantageously, data generated by the child task may be written to the queue while the TMD **322** representing the child task is executing. The queue may be implemented as a circular queue so that the total amount of data is not limited to the size of the queue.

## 11

CTAs that belong to a grid have implicit grid width, height, and depth parameters indicating the position of the respective CTA within the grid. Special registers are written during initialization in response to commands received via front end **212** from device driver **103** and do not change during execution of a processing task. The front end **212** schedules each processing task for execution. Each CTA is associated with a specific TMD **322** for concurrent execution of one or more tasks. Additionally, a single GPC **208** may execute multiple tasks concurrently.

A parameter memory (not shown) stores runtime parameters (constants) that can be read but not written by any thread within the same CTA (or any LSU **303**). In one embodiment, device driver **103** provides parameters to the parameter memory before directing SM **310** to begin execution of a task that uses these parameters. Any thread within any CTA (or any exec unit **302** within SM **310**) can access global memory through a memory interface **214**. Portions of global memory may be stored in the L1 cache **320**.

Local register file **304** is used by each thread as scratch space: each register is allocated for the exclusive use of one thread, and data in any of local register file **304** is accessible only to the thread to which the register is allocated. Local register file **304** can be implemented as a register file that is physically or logically divided into P lanes, each having some number of entries (where each entry might store, e.g., a 32-bit word). One lane is assigned to each of the N exec units **302** and P load-store units LSU **303**, and corresponding entries in different lanes can be populated with data for different threads executing the same program to facilitate SIMD execution. Different portions of the lanes can be allocated to different ones of the G concurrent thread groups, so that a given entry in the local register file **304** is accessible only to a particular thread. In one embodiment, certain entries within the local register file **304** are reserved for storing thread identifiers, implementing one of the special registers. Additionally, a uniform L1 cache **375** stores uniform or constant values for each lane of the N exec units **302** and P load-store units LSU **303**.

Shared memory **306** is accessible to threads within a single CTA; in other words, any location in shared memory **306** is accessible to any thread within the same CTA (or to any processing engine within SM **310**). Shared memory **306** can be implemented as a shared register file or shared on-chip cache memory with an interconnect that allows any processing engine to read from or write to any location in the shared memory. In other embodiments, shared state space might map onto a per-CTA region of off-chip memory, and be cached in L1 cache **320**. The parameter memory can be implemented as a designated section within the same shared register file or shared cache memory that implements shared memory **306**, or as a separate shared register file or on-chip cache memory to which the LSUs **303** have read-only access. In one embodiment, the area that implements the parameter memory is also used to store the CTA ID and task ID, as well as CTA and grid dimensions or queue position, implementing portions of the special registers. Each LSU **303** in SM **310** is coupled to a unified address mapping unit **352** that converts an address provided for load and store instructions that are specified in a unified memory space into an address in each distinct memory space. Consequently, an instruction may be used to access any of the local, shared, or global memory spaces by specifying an address in the unified memory space.

The L1 cache **320** in each SM **310** can be used to cache private per-thread local data and also per-application global data. In some embodiments, the per-CTA shared data may be

## 12

cached in the L1 cache **320**. The LSUs **303** are coupled to the shared memory **306** and the L1 cache **320** via a memory and cache interconnect **380**.

## Thread Logical Identifier

An input work buffer that includes multiple work items is assigned to a CTA for processing. The work buffer may include work items for a pixel tile, where each work item corresponds to a particular region of the pixel tile. A number of threads equal to the number of work items in the work buffer are attached to the “parent” CTA. If there is an excess of work items, then only a subset of available work items may be attached to the “parent” CTA. A logical identifier that is relative to the parent CTA is assigned to each thread and the logical identifiers increment in sequence so that the logical identifiers indicate the order in which the threads are launched. The logical identifiers are defined in a linear space rather than a multidimensional space. When needed, the threads may be configured to execute critical sections of code based on the logical identifiers. When the logical identifiers are used by each thread to index into an input work buffer, the threads of a CTA can be executed (via the logical to physical thread mapping) in the work order. Importantly, the logical identifiers for the threads are stored within the SM **310**, so even if the processing work is stored outside of the SM **310**, the next thread to process in the sequence specified by the logical identifiers can be identified. Then, the processing work for the next thread can be loaded into the SM **310** for processing.

In order to maintain a mapping from logical thread IDs to physical threads, we introduce the concept of a thread block, where each thread block corresponds to a fixed set of physical threads (e.g. 16 threads per block). When launching new work, we allocate physical threads one thread block at a time. This way, we need to maintain the mapping only at per-block granularity, instead of per-thread granularity. FIG. 4 is a conceptual diagram illustrating thread blocks of a CTA, according to one example embodiment of the present disclosure. As shown in FIG. 4, **16** CTAs may be executing within an SM **310** and each CTA can use eight different barriers each having a unique barrier identifier. A barrier may include “top” and “bottom” barrier instructions to define the start and end of a critical code section within a program. A simpler barrier includes only a “top” barrier instruction. In one embodiment, the program counter of the “top” barrier instruction is appended to a barrier identifier as a tag to allow the same barrier identifier to be used in multiple places in a single program. The program counter uniquely discriminates between different barriers when the same barrier identifier appears more than once in the program. In another embodiment, an incrementing counter is used to generate unique barrier identifiers.

Each CTA may be allocated at least one thread block, where a thread block includes 16 threads. As shown in FIG. 4, in one embodiment, the maximum number of thread blocks that may be allocated to a CTA is eight. In another embodiment, a thread block includes 64 threads and each CTA may include 512 or more threads. Sixteen warps are reserved for processing the thread blocks, where each warp includes 4 threads. Therefore, each thread block is a group of 64 threads having resources that are allocated together. As shown in FIG. 4, 128 warps may be simultaneously processed by an SM **310** and the four thread blocks may be distributed to different execution units **302** for load balancing across the different execution units **302**. In other embodi-

ments, different numbers of threads may be included in a CTA and a thread block may include a different number of threads.

Given a logical identifier associated with a particular thread, the corresponding thread group that includes the particular thread may be determined. In one embodiment, the thread block is computed by truncating the lowest 4 bits of the logical identifier. The lowest 4 bits of the logical identifier are an offset within the thread block. The physical identifier for the thread is computed by mapping the thread block to a corresponding physical identifier base and then using the offset to locate the processing resources allocated for the thread. For example, the high bits of the physical identifier may be used to determine the thread block and the lower bits may be used to determine the particular thread within the thread block. In sum, the logical identifier is a CTA-oriented identifier for each thread in the CTA and the physical identifier is hardware-oriented identifier for each thread that is executed by the SM 310 for the CTA.

#### Conditional Barrier Instruction Scheduling

FIG. 5A is a block diagram of the warp scheduler and instruction unit 312 of FIG. 3C, according to one example embodiment of the present disclosure. As shown in FIG. 5A, the warp scheduler and instruction unit 312 includes an instruction cache fetch unit 412 that is configured to fetch cache lines containing the instructions for warps from the instruction L1 cache 370. In one embodiment, each cache line is 512 bits wide, storing eight instructions (64 bits wide) in a single cache line. The instruction cache fetch unit 412 outputs the instructions to the instruction scheduling unit 420.

The instruction scheduling unit 420 receives the instructions and warp identifiers and proceeds to schedule the instructions for execution. The instruction scheduling unit 420 may be configured to maintain a priority associated with each of the warps scheduled on SM 310 and schedule the fetched instructions based on the priorities. For example, scheduling unit 425 may maintain a 6-bit or a 10-bit priority value associated with each of 16 different warps scheduled on SM 310 at any given time. The priority may be assigned based on various factors. In one embodiment, priority may be based on when the warp was scheduled on SM 310 (i.e., the longest pending warp may have the highest priority). In other embodiments, the priority may for each warp may be specified by the program that is defined by the instructions executed by the warp or may be based on the CTA.

The instruction scheduling unit 420 includes a scheduling unit 425, a thread state 510, and a barrier scheduling unit 430. The scheduling unit 425 does not necessarily select the instructions in the priority order of the different warps because one or more of the instructions may not be ready for execution due to a data dependency or because all of the threads participating in a barrier instruction have not reached a barrier instruction. When a first instruction can be issued, the instruction is scheduled and output by the scheduling unit 425. When the first instruction cannot be issued, the scheduling unit 425 determines if an instruction for a different warp may be issued for the respective instruction. In some cases, the first instruction can be issued, but the first instruction is low priority, so that another instruction (from a different warp) may be issued instead. In all cases, the instructions for each individual thread of a warp are issued in the order that the instructions for the respective individual threads are received by the warp scheduler and instruction unit 312.

The scheduling unit 425 maintains a state model of SM 310 that is updated based on the issued instructions. The state model allows the scheduling unit 425 to select instructions based on dynamic execution of the program and the availability of resources within SM 310. For example, a SM 310 or functional unit within an SM 310 or the texture unit 315 that will execute the instruction may be identified as a resource needed for the instruction and the availability of the resource may be used by the scheduling unit 425.

State information is maintained and used by the scheduling unit 425 and the barrier scheduling unit 430. In particular, state information is needed for barriers and for threads. Barrier state includes one or more of the following fields: a reference count, a rescan flag, impatient flag, a static/dynamic/idle field, and a barrier type (e.g., simple conditional, critical section, ordered critical section). A barrier begins in an idle state. When at least one thread reaches the top barrier instruction, the barrier state changes from idle to static. The scheduling unit 425 changes the state from static to dynamic when the barrier is "ready" to be executed. The barrier scheduling unit 430 changes the state from dynamic to idle when execution of the barrier is "done".

The thread state is stored in thread state 510 and includes one or more of the following fields: barrier identifier, an identifier of the parent CTA, the logical identifier, a done flag, and an awake/asleep flag. The use of the thread state 510 and the barrier state is described in further detail in the following paragraphs.

When the scheduling unit 425 identifies a first occurrence of a particular barrier instruction, meaning that a first thread of a CTA has reached the particular barrier instruction during execution of a program, the scheduling unit 425 updates the state of the barrier identifier specified by the barrier instruction from "idle" to "static" (assuming that the first thread causes an increment of the barrier membership counter that is compared with the reference count). It is not necessary for all threads of a CTA to participate in each barrier that is allocated to a CTA. Each thread that participates in a particular barrier specifies a barrier identifier corresponding to the particular barrier and a thread may participate in one barrier at a time. In one embodiment, threads may indicate participation in a particular barrier using an instruction predicate. Threads that do not participate in a barrier may be scheduled to execute instructions received by the scheduling unit 425. Threads that do participate in a barrier cannot execute instructions that are after the barrier instruction in program order until all of the participating threads have reached the barrier instruction.

The scheduling unit 425 is configured to change the barrier state from static to dynamic when the "membership" for a particular barrier is complete. A variety of different conditions may be used to determine that the membership is complete. A first condition is that the number of threads that have reached the barrier instruction equals the number of threads assigned to the parent CTA. A second condition is that the number of threads that have reached the barrier instruction equals the reference count specified for the barrier. A third condition is that the number of threads that have reached the barrier instruction and that participate in the barrier equals the reference count specified for the barrier. The reference value is specified by the program and indicates the number of threads that are expected to arrive at the barrier. A thread that does not participate in a barrier stores zero for the barrier identifier in the thread state 510. The scheduling unit 425 may be configured to increment a barrier membership counter corresponding to the specific

15

barrier as threads reach the barrier instruction to determine if the membership for the barrier is complete.

As each participating thread reaches the barrier instruction, state data for the thread that is stored in the thread state 510 is updated. Specifically, the thread state is set to “asleep,” indicating that the thread is halted and is not executing instructions. Once the membership is complete, the scheduling unit 425 changes the state of the barrier identifier to “dynamic” and outputs the CTA identifier and the barrier identifier to the barrier scheduling unit 430. The barrier scheduling unit 430 is configured to schedule threads participating in barrier instructions for execution.

Barrier instructions may be used to delineate ordered and non-ordered critical code sections of a program. A top barrier instruction occurs immediately prior to the first instruction of a critical code section and a bottom barrier instruction having the same barrier identifier as the top barrier instruction occurs immediately after the last instruction of the critical code section. TABLE 1 illustrates an example of an ordered critical code section.

TABLE 1

|                      |   |
|----------------------|---|
| BARRIER.TOP.OCS      | // start ordered critical section           |
| LB R0, [address];    | // load CTA counter into R0, for local work |
| IMAD R2, R0, R1, R3; | // increment using thread values            |
| ST [address], R2     | // store the CTA counter from R2 to memory  |
| BARRIER.BOT.OCS;     | // finish the ordered critical section      |

Barrier instructions may also be used to delineate critical code sections that do not have order constraints, i.e., non-ordered critical code sections. Non-ordered critical code sections may be scheduled by the barrier scheduling unit 430 in an arbitrary order or the logical identifiers may be used to schedule the critical code sections in the logical identifier order, just as ordered critical code sections are scheduled. The barrier scheduling unit 430 will only schedule the (ordered and non-ordered) critical code sections for execution by the threads that participate in the barrier. In one embodiment, a barrier instruction does not delineate a critical code section and is used as a (simple) conditional barrier to synchronize execution of participating threads at a barrier instruction. In other words, instead of a top and bottom barrier, there is just a single barrier instruction.

In some cases, impatient barrier scheduling may be enabled and the scheduling unit 425 is configured to transition the barrier into the dynamic state without necessarily waiting for all of the threads that participate in the barrier to reach the barrier instruction. In other words, a subset of the participating threads can start execution of the barrier before at least one other participating thread has reached the barrier. Each barrier may specify whether impatient scheduling is enabled for the barrier. When impatient barrier scheduling is enabled, additional conditions may be used to determine that the membership is complete.

As previously explained, the scheduling unit 425 is configured to use a variety of different conditions to determine that the membership for a particular barrier is complete. A fourth condition is that a timeout occurs. The timeout may be measured in a variety of different ways. The timeout may be measured from the first (participating) thread to reach the barrier. The timeout may be measured from the first (participating or non-participating) thread to reach the barrier. The timeout may be measured from the most recent (participating) thread to reach the barrier. The timeout may be measured from the most recent (participating or non-par-

16

ticipating) thread to reach the barrier. A fifth condition is that either a timeout occurs or a minimum number of threads have reached the barrier.

An impatient barrier functions as a “speed bump” to encourage grouping of some threads while not requiring the threads to wait for all threads that participate in the barrier to reach the barrier. Therefore, at least a portion of threads expected to have localized memory access patterns may be grouped for improved cache access performance. An impatient barrier may be used to execute code that is non-ordered critical. An impatient barrier may be used to execute code that is ordered critical, however, the threads participating in the impatient barrier may not necessarily be executed in order because an earlier (in terms of order) thread may reach the barrier after a later (in terms of order) thread has already been executed.

When a barrier is scheduled as an impatient barrier (i.e., only a portion of the participating threads have reached the barrier), the scheduling unit 425 may be configured to invalidate the barrier so that any participating threads that arrive after the barrier is transitioned into the dynamic state ignore the barrier and do not wait for the other participating threads. In general, when the barrier is used to delineate a critical section of code, the barrier should not be invalidated. Invalidation should only be performed on a conditional barrier that does not delineate a critical code section to allow late arriving threads to continue execution, as described further herein. If threads arrive late to an invalidated critical section of code, the late arriving threads should not be allowed to execute (e.g., via predication) the delineated code to prevent destructive interference.

#### Late Arriving Thread Handling

When impatient scheduling is enabled for a barrier, threads that arrive after a barrier has transitioned from static to dynamic may or may not be allowed to “join” the barrier late. Whether or not late arriving threads may join an impatient barrier may be specified for each barrier or may be enabled for all barriers that can be scheduled as impatient barriers. The done flag is maintained so that the barrier scheduling unit 430 may differentiate between a late arriving thread that has not yet been scheduled (done=0 and asleep=1) and a thread (late arriving or not) that has already executed the critical code section (done=1), but that is after the late arriving thread in logical order. The done flag is cleared when each participating thread reaches the barrier instruction and the done flag is set when the thread executes the bottom barrier instruction that delineates the critical code section.

When a late arriving participating thread joins an impatient barrier, the scheduling unit 425 sets the thread state to done=0 and asleep=1, indicating that the late arriving participating thread is halted and is not executing instructions. The scheduling unit 425 also sets the rescan flag for the barrier. The rescan flag indicates that a late arriving participating thread has joined the impatient barrier. When the rescan flag is set, the barrier scheduling unit 430 will rescan the thread state 510 to identify and schedule any late arriving participating threads, i.e., threads that do not have their thread state set as done=1. The rescan would optimally be performed at the end of execution of the current scan, but may be performed at any time up to the end of execution of the current scan. If the impatient barrier is not accepting late arriving threads, the thread will be prevented from issue by

barrier scoreboarding logic (not shown). Once the barrier transitions back to idle, the thread is allowed to issue and start the barrier anew.

In one embodiment, the scheduling unit 425 is configured to allow any late arriving participating threads to ignore the barrier after waiting for the portion of participating threads that reached the barrier to execute. The scheduling unit 425 keeps the late arriving participating threads “asleep” (with cleared done flags) until after execution of the portion of participating threads that reached the barrier is complete. When the portion of participating threads complete execution of the barrier, the scheduling unit 425 updates the thread state 510 for the late arriving participating threads to “awake”, and the done flags are set for all of the participating threads. The late arriving participating threads and the threads in the portion of participating threads then continue execution. The scheduling unit 425 may configure the late arriving threads to execute the code delineated by the barrier or to resume execution at the same point as the threads that completed execution of the barrier.

#### Barrier Scheduling

The barrier scheduling unit 430 schedules the threads that participate in a barrier by waking up a first thread, where the first thread is the participating thread having the lowest logical identifier value (unless a different ordering convention is used). The barrier scheduling unit 430 updates the thread state 510 to indicate that the first thread is awake. Consequently, the scheduling unit 425 will schedule the first thread for execution because the first thread is now eligible to be executed. When the first thread reaches the bottom barrier, the barrier scheduling unit 430 is notified by the scheduling unit 425 and the thread state 510 is updated by either the barrier scheduling unit 430 or the scheduling unit 425 to indicate that the first thread is asleep. When the first thread reaches the bottom barrier, the done flag in the thread state 510 is also updated by either the barrier scheduling unit 430 or the scheduling unit 425 to indicate that the first thread is done.

The barrier scheduling unit 430 then wakes up the next participating thread, enabling execution of the critical code section for the next participating thread. The barrier scheduling unit 430 continues to wake each participating thread in logical order until the last participating thread reaches the bottom barrier. When the last participating thread reaches the bottom barrier, the barrier scheduling unit 430 may be configured to check for late arriving threads, and, after confirming that all participating threads have done flags that are set, execution of the barrier is complete and the barrier scheduling unit 430 updates the thread state 510 to indicate that all of the participating threads are awake. The barrier scheduling unit 430 may determine that the last participating thread has been scheduled when the search for a next participating thread fails. The scheduling unit 425 updates the barrier state to indicate that the barrier identifier is “idle”, e.g., neither “static” nor “dynamic”.

When at least one late arriving participating thread has joined an impatient barrier, the rescan flag is set in the barrier state. The barrier scheduling unit 430 may be configured to search for late arriving threads when waking up the next participating thread. Because a late arriving thread may be “out-of-order” in terms of logical order, the barrier scheduling unit 430 may awaken participating threads in an order that is not strictly the logical order. When the last participating thread reaches the bottom barrier and the rescan flag is set, the barrier scheduling unit 430 should

check for late arriving participating threads before updating the barrier identifier state to indicate that the barrier identifier is idle. The barrier scheduling unit 430 clears the rescan flag whenever the barrier scheduling unit 430 completes a rescan to search for late arriving threads.

When the barrier is an impatient conditional barrier and the barrier identifier state is changed from static to dynamic, the scheduling unit 425 may be configured to invalidate the barrier identifier so that no late arriving participating threads can join the impatient barrier. Only impatient conditional barriers that were scheduled without waiting for all participating threads may be invalidated. In one embodiment, the scheduling unit 425 updates the count that is maintained for the barrier identifier of a conditional barrier when each late arriving thread joins an impatient barrier and the barrier identifier is only invalidated when the count does not equal the reference value and/or the barrier timeout occurs, when the barrier identifier is transitioned from dynamic to idle. In still another embodiment, the barrier identifier is not invalidated and if any other participating threads reach the barrier, the count is updated and the scheduling unit 425 transitions the barrier from the static to dynamic state when the count equals the reference value or the barrier is processed again as an impatient barrier to schedule another portion of the participating threads.

The barrier scheduling unit 430 may also be configured to schedule exclusive critical code sections and ordered critical code sections with a non-blocking bottom barrier. The barrier scheduling unit 430 schedules a first thread participating in an exclusive critical code section by first waiting for execution of any other critical code sections or exclusive critical code sections to complete. The barrier for an exclusive critical code section is scheduled for exclusive execution. An exclusive critical code section may be an ordered critical code section or a non-ordered critical code section that is exclusive. Note, that threads that are not participating in a barrier may be executed simultaneously with the threads executing the exclusive critical code section. An exclusive critical code section may be used when resource access conflicts may occur between the threads participating in different barriers.

The barrier scheduling unit 430 schedules threads participating in a critical code section with a non-blocking bottom barrier by allowing the threads to continue execution of subsequent instructions that are immediately after the bottom barrier without waiting for all of the participating threads to execute the critical code section. Once the thread state 510 for a participating thread is updated by the barrier scheduling unit 430 to indicate that the thread is awake, the thread remains awake after reaching the bottom barrier. The scheduling unit 425 is then able to schedule participating threads that have completed execution of the critical code section for execution simultaneously with one other participating thread that is currently executing the critical code section. When all of the participating threads have reached the bottom barrier, execution of the barrier is complete and the barrier scheduling unit 430 updates the barrier state to indicate that the barrier identifier is neither “static” nor “dynamic”.

In one embodiment, the scheduling unit 425 performs a dynamic barrier scoreboard check to determine whether a first instruction can be issued before outputting the first instruction to the decode unit 450. When the first instruction is associated with a barrier identifier, the scheduling unit 425 waits to issue the first instruction in the following situations: when a barrier (in a static or dynamic state) has the same barrier identifier and a different tag compared with the first

instruction or when a barrier (in the dynamic state) has the same identifier and the same tag and late arriving threads are not accepted for the barrier identifier.

Decode unit 450 receives the next instruction to be dispatched from the instruction scheduling unit 420. The decode unit 450 performs a full decode of the instruction and transmits the decoded instruction to the dispatch unit 470. Again, in some embodiments, instructions may be dual or quad issued and decode unit 450 may implement separate decode logic for each issued instruction. Dispatch unit 470 implements a FIFO and writes the decoded values to local register file 304 for execution by execution units 302 or load/store units 303. In embodiments that issue multiple instructions simultaneously, dispatch unit 470 may issue each instruction to a different portion of the functional units of SM 310. Scoreboard unit 480 manages and tracks the number of instructions that have been decoded and dispatched per thread group.

FIG. 5B is a block diagram of a portion of the scheduling unit 425 and the barrier scheduling unit 430 of FIG. 5A, according to one example embodiment of the present disclosure. A thread barrier arrival tracking unit 500 updates the barrier state information for each barrier that is allocated to a CTA. As previously explained, the barrier state information indicates which barrier identifiers are "static", "dynamic", or idle (neither "static" nor "dynamic") and includes a rescan flag for each barrier identifier that is set when at least one late arriving participating thread has joined an impatient barrier. When a barrier identifier transitions from "static" to "dynamic", the barrier identifier and CTA identifier is output to the FIFO 505. As each thread of the CTA that participates in the barrier reaches the barrier instruction, the scheduling unit 425 updates the state of the participating thread that is stored in the thread state 510 to indicate that the participating thread is "asleep".

The FIFO 505 buffers dynamic barriers when an execution slot is not available in the barrier execution slots 515. In one embodiment, the barrier execution slots 515 includes 16 slots that may each be occupied by one dynamic barrier. When a slot is available in the barrier execution slots 515, a dynamic barrier is popped from the FIFO 505 and inserted into the slot. The barrier arbitration unit 520 arbitrates between the different dynamic barriers that occupy the slots of the barrier execution slots 515. Various priority schemes may be employed by the barrier arbitration unit 520 to arbitrate between the different dynamic barriers. The barrier arbitration unit 520 provides the thread selection unit 530 with a dynamic barrier from which a thread may be selected for scheduling by the thread selection unit 530.

The thread selection unit 530 selects threads in the order specified by the logical identifier associated with each thread in a thread array (i.e., allocated to a CTA). The thread selection unit 530 accesses the barrier participation tracking information 532 to determine which threads participate in the dynamic barrier. The thread selection unit 530 accesses the thread state 510 and the barrier state information (not shown) to identify any late arriving participating threads.

Not all threads in the one or more thread blocks that are allocated to a CTA necessarily participate in each barrier that is used by the CTA. As previously explained, the barriers are specified by barrier identifiers and each thread indicates whether or not it participates in one of the barriers. The thread selection unit 530 identifies the participating threads once during processing of the barrier and then proceeds to select each participating thread for execution serially. The thread selection unit 530 skips over non-participating threads during the selection process. When a barrier is

impatient, during the selection process the thread selection unit 530 skips over non-participating threads and threads that have not yet arrived at the barrier (or that did not arrive when the barrier became dynamic if late arriving threads cannot join the barrier).

Before selecting a first thread for execution for a particular barrier, the thread selection unit 530 determines if the barrier is an exclusive barrier that delineates an exclusive critical code section. If the barrier is exclusive, then the thread selection unit 530 determines if any other thread arrays are executing a critical code section or exclusive critical code section, and, if so, the thread selection unit 530 waits until threads in those thread arrays complete execution before selecting a first thread for execution for the exclusive barrier. In one embodiment, the FIFO 505 does not pop an exclusive barrier while a dynamic barrier occupies a slot in the barrier execution slots 515 and the FIFO 505 does not pop a barrier while an exclusive barrier occupies a slot in the barrier execution slots 515.

The thread selection unit 530 may generate an execution mask based on the participating threads. Participating threads specify the barrier identifier that matches the barrier identifier provided with the barrier instruction. When a late arriving thread is allowed to join a barrier, the thread selection unit 530 may regenerate the execution mask, setting bits in the mask for participating threads having a negated "done" flag. The thread selection unit 530 searches the execution mask to find the first thread, in logical identifier order, to select for execution until all of the threads have been selected. As each participating thread is selected, the bit of the execution mask corresponding to the participating thread is cleared and the thread is marked as done. In one embodiment, when multiple thread blocks are allocated to a CTA, the thread selection unit 530 generates an execution mask for one thread block at a time, thereby limiting the number of bits in the execution mask to the number of threads in a thread block.

When the thread selection unit 530 selects a participating thread for execution, the thread selection unit 530 updates the state of the thread stored in the thread state 510 to indicate that the thread is "awake". The scheduling unit 425 will then issue the thread for execution and continue to issue the thread for each instruction in the critical code section until the barrier bottom is reached. When the bottom barrier is reached, the scheduling unit 425 informs the thread selection unit 530 and the thread selection unit 530 determines whether the thread is required by the barrier to wait for all other threads participating in the barrier to execute the critical code section before proceeding to execute an instruction that is after the critical code section in program order, i.e., the thread selection unit 530 determines if the bottom barrier is a non-blocking bottom barrier. When the barrier is a non-blocking barrier, the thread selection unit 530 may select a next participating thread for execution without putting the currently selected participating thread to "sleep". Instead, the state of the next participating thread that is stored in the thread state 510 is updated to "awake" and any participating threads that have completed execution of the critical code section continue to execute subsequent instructions of the program in program order. The thread selection unit 530 also sets the respective done flag as each participating thread reaches the bottom barrier.

In one embodiment, barriers are used to ensure that threads for a particular CTA that processes data using the texture unit 315 are performed without allowing threads from a different CTA to intervene. Ensuring that the threads for one CTA are processed by the texture unit 315 together

21

increases the likelihood of cache hits because the texture accesses are localized within a CTA. The texture unit 315 is an example of a shared resource and the barriers may be used to control which threads access a shared resource or a resource that may benefit from locality. While the threads may not need to be executed in a specific order, delineating the texture load instructions which read texture data from memory as exclusive critical code sections provides a mechanism for increasing the likelihood that the texture reads will hit in the L1.5 cache 335. Instructions in exclusive critical code sections that control a specific resource (e.g. the texture unit 315) are able to execute at the same time as critical code sections and/or exclusive critical code sections that do not use the same resource. The barrier execution slots 515 may include a slot dedicated to barrier delineating instructions that are executed by the texture unit 315, or slots for any resources that prefer localized controlled access. Threads that are executed by the texture unit 315 may be scheduled for execution simultaneously with threads that are executed by the execution units 302. Scheduling priority may be specified for the different barrier identifiers and texture load instructions delineated using a barrier instruction may or may not be scheduled with higher priority than other barriers.

FIG. 6A illustrates a flowchart 600 of a method for scheduling instructions for execution based on logical identifiers, according to one example embodiment of the present disclosure. Although the method steps are described in conjunction with the systems of FIGS. 1, 2, 3A-3C, 5A, and 5B persons of ordinary skill in the art will understand that any system configured to perform the method steps, in any order, is within the scope of the disclosure.

The method 600 begins at step 605, where the threads allocated to a CTA are associated with logical identifiers. The logical identifiers are mapped to physical identifiers that are referenced by the SMs 310 during execution of the threads. At step 610, program execution is initiated and the scheduling unit 425 receives instructions. At step 615, the scheduling unit 425 determines if an instruction is a barrier instruction, and, if not, at step 658, the scheduling unit 425 continues execution of the program by outputting the instruction to the decode unit 450.

If, at step 615, the scheduling unit 425 determines that the instruction is a barrier instruction, and, if so, then at step 616 the scheduling unit 425 updates the barrier membership according to the membership condition that is used for the barrier. When the barrier instruction is reached for the first thread of a CTA, the scheduling unit 425 also updates the state of the barrier to indicate that the barrier is "static" (assuming that the thread causes an increment of the barrier membership counter that is compared with the reference count). Then, at step 617, the scheduling unit 425 determines if the thread participates in the barrier instruction. If the thread does not participate in the barrier instruction, then the scheduling unit 425 proceeds to step 658, and the thread continues execution of the program. Otherwise, the thread participates in the barrier instruction, and, at step 620, the scheduling unit 425 updates the state of the thread as "asleep". When the barrier instruction delineates the start of a non-impatient critical code section (ordered or non-ordered), the scheduling unit 425 membership condition should require all of the participating threads in the CTA to reach the barrier instruction before allowing any of the participating threads to execute the critical code section.

When the barrier is an impatient barrier, the scheduling unit 425 need not wait for all participating threads to reach the barrier instruction before allowing a portion of the

22

participating threads to be scheduled for execution. Therefore, at step 621, the scheduling unit 425 determines if the thread participates in an impatient barrier whose state has already transitioned from static to dynamic and if late arriving threads are allowed to join the barrier. In other words, the scheduling unit 425 determines if the thread is a late arriving thread that may join an impatient barrier, and if the thread is a late arriving thread that may join an impatient barrier, at step 625, the scheduling unit 425 sets the rescan flag for the impatient barrier. The done flag will have already been cleared for the late arriving thread. The barrier scheduling unit 430 will then detect and be able to select the late arriving thread for scheduling. If, at step 621, the scheduling unit 425 determines that the thread is a late arriving thread that may not join an impatient barrier, then the scheduling unit 425 proceeds to step 622. The scheduling unit 425 will also proceed to step 622 when the barrier is not an impatient barrier. In one embodiment, scoreboarding of barriers may be used to indicate whether an impatient barrier is accepting late arriving threads.

At step 622, the scheduling unit 425 determines if the membership is complete, according to the variety of different conditions that may be used to determine that the membership is complete, and, if not, at step 658, the participating threads that have reached the barrier remain asleep while the scheduling unit 425 continues execution of the program. Otherwise, the scheduling unit 425 updates the state of the barrier from "static" to "dynamic" and outputs the barrier identifier and the CTA identifier of the thread to the barrier scheduling unit 430 before proceeding to step 655.

At step 655, the barrier scheduling unit 430 selects participating threads for execution. When all of the participating threads have completed execution of the critical code section, program execution continues at step 658. Additional details of step 655 are described in conjunction with FIG. 6B. In one embodiment, if a barrier becomes "hung" so that the threads cannot be executed, the barrier can be reset just the same way that invalidated barriers can be reset, i.e., via a special instruction.

In one embodiment, a conditional barrier is used to synchronize execution of participating threads at a barrier instruction, and, at step 620, the thread state for each participating thread is set to asleep and when the membership is complete at step 622, the scheduling unit 425 updates the state of the barrier from "static" to "dynamic" and updates the state of all of the participating threads stored in the thread state 510 to "awake". At step 655, the participating threads (or a portion of the participating threads when the conditional barrier is scheduled impatiently) then continue execution of the program and the barrier scheduling unit 430 updates the state of the barrier to neither "static" nor "dynamic", i.e., indicating that execution of the barrier is complete for at least a portion of the threads.

FIG. 6B illustrates a method for performing step 655 shown in FIG. 6A, according to one example embodiment of the present disclosure. Although the method steps are described in conjunction with the systems of FIGS. 1, 2, 3A-3C, 5A and 5B persons of ordinary skill in the art will understand that any system configured to perform the method steps, in any order, is within the scope of the disclosure.

At step 630, the barrier scheduling unit 430 selects a thread for execution based on the logical identifiers associated with the threads that participate in the barrier. The state of the selected thread is updated from "asleep" to "awake" so that the thread is eligible for execution. Additional details

23

of step 630 are described in conjunction with FIG. 6C. When the rescan flag is set, the barrier scheduling unit 430 detects any late arriving threads and may select threads for execution out-of-sequence compared with the logical identifiers. More specifically, when the barrier scheduling unit 430 schedules the “last” thread and the rescan flag is set, the barrier scheduling unit 430 rescans the execution mask, ignoring threads having a “done” flag that is set, and schedules any late arriving threads.

At step 635, the instructions within the critical code section are dispatched by the dispatch unit 470 and are executed for the selected thread. At step 636, the scheduling unit 312 receives the bottom barrier instruction associated with the barrier identifier and the state of the thread is updated to set the done flag. At step 638, the scheduling unit 312 determines if the bottom barrier instruction is a blocking bottom barrier instruction. When, at step 638, the scheduling unit 312 determines that the bottom barrier instruction is blocking, then, at step 640, the barrier scheduling unit 430 determines if another thread that participates in the barrier needs to execute the critical code section. If another thread needs to execute the critical code section, then at step 643, the barrier scheduling unit 430 updates the state of the thread that has reached the bottom barrier instruction is updated from “awake” to “asleep” so that the thread is not eligible for execution. Otherwise, at step 650, the barrier scheduling unit 430 updates the thread state 510 to indicate that all threads participating in the barrier are “awake” before proceeding to step 652.

If, at step 638 the bottom barrier instruction is not a blocking bottom barrier instruction, then at step 642, the barrier scheduling unit 430 determines if another thread that participates in the barrier needs to execute the critical code section. If another thread needs to execute the critical code section, the barrier scheduling unit 430 returns to step 630. Otherwise, at step 652, the barrier scheduling unit 430 removes the barrier and CTA identifiers from an execution slot in the barrier execution slots 515 and updates the state of the barrier to indicate that the barrier is neither “static” nor “dynamic”. At this point, when the barrier is an impatient barrier, the scheduling unit 425 may invalidate the barrier identifier if the barrier does not delineate a critical code section (i.e., the barrier is a conditional barrier) so that late arriving threads can no longer join the barrier.

The scheduling unit 425 then proceeds to continue execution of the program. During execution of a critical code section for one CTA barrier, other threads of other barriers of the same CTA, as well as thread barriers of other CTAs, may also continue execution of other code. However, if the barrier is exclusive, only program instructions that are not included within critical code sections may be executed for other barriers.

FIG. 6C illustrates a method for performing step 630 shown in FIG. 6A, according to one example embodiment of the present disclosure. Although the method steps are described in conjunction with the systems of FIGS. 1, 2, 3A-3C, 5A and 5B persons of ordinary skill in the art will understand that any system configured to perform the method steps, in any order, is within the scope of the disclosure.

The step 630, is performed by the barrier scheduling unit 430 for each scheduling cycle. At step 660, the barrier scheduling unit 430 identifies any pending barriers in the barrier execution slots 515 storing a dynamic barrier and having a next thread that can be scheduled. If an exclusive barrier is being executed, then the barrier scheduling unit 430 waits until execution of the exclusive barrier is complete

24

before identifying any pending barriers. At step 662, the barrier scheduling unit 430 selects the highest priority dynamic barrier. At step 665, the barrier scheduling unit 430 determines if the selected dynamic barrier is an exclusive barrier, and, if so, at step 668, the barrier scheduling unit 430 waits until any other barriers delineating (ordered or non-ordered) critical code sections are not being executed by threads before proceeding to step 670. Note that the barrier scheduling unit 430 may wait for all participating threads to complete execution of the other critical code section or only wait for threads currently executing to complete execution of the critical code section before (serially) executing the threads participating in the exclusive barrier. Note that multiple non-exclusive barriers may be executed in parallel for multiple CTAs, but only one exclusive barrier may be executed at a time.

At step 670, the barrier scheduling unit 430 identifies the threads that participate in the selected barrier by determining whether each thread specifies the barrier by name. At step 675, the barrier scheduling unit 430 searches for the next participating thread by examining the logical identifiers associated with each participating thread that has not yet executed the critical code section. At step 680, the barrier scheduling unit 430 updates the thread state stored in the thread state 510 for the selected participating thread to indicate that the thread is “awake” before proceeding to step 635. Note that when the rescan flag is not set, step 670 may be performed once, once per thread block, or once per cycle, to identify the participating threads because the barrier scheduling unit 430 may be configured to maintain an execution mask. The barrier scheduling unit 430 efficiently wakes up the participating threads in logical order as each previous participating thread completes execution of the critical code section and updates the execution mask. In one embodiment, the execution mask is computed for each thread block. When the rescan flag is set, step 670 may be performed once per cycle to identify the participating threads that have reached the barrier.

While various embodiments have been described above, it should be understood that they have been presented by way of example only, and not limitation. Thus, the breadth and scope of a preferred embodiment should not be limited by any of the above-described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents.

What is claimed is:

1. A method comprising:

initiating execution of a plurality of threads to process instructions of a program that includes a barrier instruction;

for each thread in the plurality of threads, determining whether the thread participates in the barrier instruction when the thread reaches the barrier instruction during execution of the thread;

prior to executing each of the threads that participate in the barrier instruction, waiting for a portion of the threads that participate in the barrier instruction to reach the barrier instruction and for conditions that allow the barrier instruction to be scheduled as an first barrier to be met; and

serially executing each of the threads that participate in the barrier instruction to process one or more instructions of the program that follow the barrier instruction, wherein during execution of each of the threads that participate in the barrier instruction an additional thread is identified as a late arriving thread that participates in the barrier instruction and that is not included in the

25

portion of the threads and the additional thread is executed to process the one or more instructions of the program that follow the barrier instruction.

2. The method of claim 1, further comprising, prior to executing each of the threads that participate in the barrier instruction, waiting for a last arriving thread that participates in the barrier instruction, wherein the last arriving thread reaches the barrier instruction last relative to other threads in the plurality of threads that participate in the barrier instruction.

3. The method of claim 1, further comprising, after executing each of the threads in the portion of the threads that participate in the barrier instruction, invalidating a barrier identifier that corresponds to the barrier instruction.

4. The method of claim 1, wherein the conditions that allow the barrier instruction to be scheduled as the first barrier comprise the portion of the threads that participate in the barrier instruction including a number of threads that is greater than a minimum number of threads.

5. The method of claim 1, wherein the conditions that allow the barrier instruction to be scheduled as the first barrier comprise a timeout expiring.

6. The method of claim 1, wherein the barrier instruction delineates a code section.

7. The method of claim 6, further comprising, waiting for a first thread to execute all instructions within the code section before selecting another thread to execute the code section.

8. The method of claim 1, further comprising associating the threads with logical identifiers that are mapped to physical identifiers, wherein the physical identifiers are referenced by a multi-threaded processing core during execution of the threads.

9. The method of claim 8, wherein executing each of the threads that participate in the barrier instruction comprises selecting the threads for serial execution based on the logical identifiers.

10. The method of claim 1, further comprising appending a tag to a barrier identifier corresponding to the barrier instruction to uniquely identify each occurrence of the barrier instruction when the barrier identifier is used in multiple places in the program.

11. A processing subsystem comprising:

an instruction scheduling unit, configured to:

receive instructions of a program for execution by a plurality of threads wherein the program includes a barrier instruction;

when each thread reaches the barrier instruction during execution of the thread, determine whether the thread participates in the barrier instruction; and

26

wait for a portion of the threads that participate in the barrier instruction to reach the barrier instruction, wherein the portion comprises a number of threads that is greater than or equal to a minimum number of threads;

after waiting for the portion of the threads, enable the threads that participate in the barrier instruction to be selected for execution; and

a multi-threaded processing core that is configured to serially execute each of the threads that participate in the barrier instruction to process one or more instructions of the program that follow the barrier instruction, wherein the instruction scheduling unit is further configured to, during the execution of the enabled threads that participate in the barrier instruction identify an additional thread as a late arriving thread that participates in the barrier instruction and enable the additional thread to be selected for execution.

12. The processing subsystem of claim 11, wherein the instruction scheduling unit is further configured to, prior enabling each of the threads that participate in the barrier instruction, wait for a last arriving thread that participates in the barrier instruction, wherein the last arriving thread reaches the barrier instruction last relative to other threads in the plurality of threads that participate in the barrier instruction.

13. The processing subsystem of claim 11, wherein the instruction scheduling unit is further configured to, after executing each thread in the portion of the threads that participate in the barrier instruction, invalidate a barrier identifier that corresponds to the barrier instruction.

14. The processing subsystem of claim 11, wherein the barrier instruction delineates a code section.

15. The processing subsystem of claim 14, wherein the instruction scheduling unit is further configured to, wait for a first thread to execute all instructions within the code section before selecting another thread to execute the code section.

16. The processing subsystem of claim 11, wherein the threads are associated with logical identifiers that are mapped to physical identifiers, and wherein the physical identifiers are referenced by the multi-threaded processing core during execution of the threads that participate in the barrier instruction.

17. The processing subsystem of claim 16, wherein the instruction scheduling unit is further configured to select the threads that participate in the barrier instruction for serial execution based on the logical identifiers.

\* \* \* \* \*